When Are Welfare Guarantees Robust?*†

Tim Roughgarden[1], Inbal Talgam-Cohen[2], and Jan Vondrák[3]

1 **Stanford University, Stanford, CA, USA**
`tim@cs.stanford.edu`
2 **Hebrew University of Jerusalem, Jerusalem, Israel**
`inbal.talgamcohen@mail.huji.ac.il`
3 **Stanford University, Stanford, CA, USA**
`jvondrak@stanford.edu`

Abstract

Computational and economic results suggest that social welfare maximization and combinatorial auction design are much easier when bidders' valuations satisfy the "gross substitutes" condition. The goal of this paper is to evaluate rigorously the folklore belief that the main take-aways from these results remain valid in settings where the gross substitutes condition holds only approximately. We show that for valuations that pointwise approximate a gross substitutes valuation (in fact even a linear valuation), optimal social welfare cannot be approximated to within a subpolynomial factor and demand oracles cannot be simulated using a subexponential number of value queries. We then provide several positive results by imposing additional structure on the valuations (beyond gross substitutes), using a more stringent notion of approximation, and/or using more powerful oracle access to the valuations. For example, we prove that the performance of the greedy algorithm degrades gracefully for near-linear valuations with approximately decreasing marginal values; that with demand queries, approximate welfare guarantees for XOS valuations degrade gracefully for valuations that are pointwise close to XOS; and that the performance of the Kelso-Crawford auction degrades gracefully for valuations that are close to various subclasses of gross substitutes valuations.

1998 ACM Subject Classification F.2 Analysis of Algorithms and Problem Complexity

Keywords and phrases Valuation (set) functions, gross substitutes, linearity, approximation

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2017.22

1 Introduction

Welfare maximization in combinatorial auctions is a central problem in both theory and practice, and is perhaps the most well-studied problem in algorithmic game theory (e.g., [7]). The problem is: given a set M of distinct items and descriptions of, or oracle access to, the valuation functions $v_1, \ldots, v_n$ of n bidders (each a set function from bundles to values), determine the partition $S_1, \ldots, S_n$ of items that maximizes the social welfare $\sum_{i=1}^{n} v_i(S_i)$.

This paper focuses on the purely algorithmic and approximation aspects of welfare maximization. Many possibility and impossibility results for efficiently computing or approximating the maximum social welfare are known, as a function of the set of allowable bidder valuations. Very roughly, the current state of affairs can be summarized by a trichotomy:

* A full version of the paper is available at `https://arxiv.org/abs/1608.02402`.

† This work is supported by NSF grants CCF-1215965 and CCF-1524062, and has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Sklodowska-Curie grant agreement No. 708935.

Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2017).
Editors: Klaus Jansen, José D. P. Rolim, David Williamson, and Santosh S. Vempala; Article No. 22; pp. 22:1–22:23
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

(i) when the valuations satisfy "gross substitutes," then exact welfare maximization is easy;
(ii) when the valuations are "complement-free" but not necessarily gross substitutes (e.g., subadditive), exact welfare maximization is hard but constant-factor approximations are possible; and
(iii) with sufficiently general valuations, even approximate welfare maximization is hard.[1]

The rule of thumb that "substitutes are easy, complements are hard" has its origins in the economic literature on multi-item auction design. For example, simple and natural ascending auctions converge to a welfare-maximizing Walrasian equilibrium (i.e., to a market-clearing price vector and allocation) whenever all bidders' valuations satisfy gross substitutes [24], but when this condition is violated a Walrasian equilibrium does not generally exist [18, 28]. Similarly, the VCG mechanism has a number of desirable properties (like revenue monotonicity) when bidders' valuations satisfy gross substitutes, but not in general otherwise [2].

Thus both computational and economic results suggest that social welfare maximization and combinatorial auction design are much easier when bidders' valuations satisfy the gross substitutes condition. More generally, in settings with both substitutes and complements, the folklore belief in the field is that simple auction formats should produce allocations with near-optimal social welfare if and only if the substitutes component is in some sense "dominant."[2] For over twenty years, there has been a healthy (and high-stakes) debate over whether or not the ideal case of gross substitutes valuations should guide combinatorial auction design in realistic settings.[3]

The goal of this paper is to evaluate rigorously the folklore belief that the main takeaways from the study of gross substitutes valuations are robust, i.e., remain valid in settings where the substitutes condition holds only approximately. We are interested both in possibility/impossibility results for achieving robustness (Section 3), and also in understanding the robustness of standard algorithms and auctions for welfare maximization (Sections 4–5, Appendix D). Our goal relates to a wider research agenda on the robustness or stability of "nice" classes of valuations: For different notions of closeness to nice valuations (in our case gross substitutes or even simply linear valuations), do fundamental optimization tasks (in our case welfare maximization) maintain their algorithmic guarantees? We consider a standard notion of pointwise closeness (Section 3), and subsequently add to it "semantic" closeness (Section 4).[4] This research agenda is motivated by inherent mathematical interest in classes of set functions and their stability [10, 14], as well as by applications like data-driven optimization in which parameters of the optimization problem such as valuations are approximately estimated from data [5, 4, 43, 19, 3].

1.1 Our Results

We first consider arguably the most natural notion of "approximate gross substitute valuations," namely valuations that are pointwise within a $1+\epsilon$ factor of some gross substitutes

[1] "Gross substitutes" means that a bidder's demand for an item can only increase as prices of other items increase; see Section 2.1 for a formal definition.
[2] For example, Bykowsky et al. [8] write: "In general, synergies across license valuations complicate the auction design process. Theory suggests that a 'simple' (i.e., non-combinatorial) auction will have difficulty in assigning licenses efficiently in such an environment."
[3] Ausubel et al. [1] write: "A contentious issue in the design of the Federal Communications Commission (FCC) auctions of personal communications services (PCS) licenses concerned the importance of synergies. If large synergies are prevalent among the licenses being offered, then the simultaneous ascending auction mechanism the FCC adopted, which does not permit all-or-nothing bids on sets of licenses, might be expected to perform poorly."
[4] What does it *mean* for a valuation to belong to a nice class? If a valuation approximately maintains the meaningful properties of the class, we say it is *semantically* close to it.

valuation. Do the laudable properties of gross substitutes valuations degrade gracefully as ϵ increases? At this level of generality, the answer is negative: even for valuations that are pointwise close to *linear* (affine) valuations, we prove that the social welfare cannot be approximated to within a subpolynomial factor using a subexponential number of value queries, and that demand oracles cannot be approximated in any useful sense by a subexponential number of value queries. When the gross substitutes condition holds exactly, a demand query can be implemented using a polynomial number of value queries (see [36]), and the welfare maximization problem can be solved exactly with a polynomial number of such queries (see [33]). We conclude that there is no sweeping generalization of the properties of gross substitutes valuations to approximations of such valuations, and that any positive result must impose additional structure on the valuations (beyond gross substitutes), use a more stringent notion of approximation, and/or use more powerful oracle access to the valuations.

We next consider positive results in the value oracle model. Our main result here is a proof that the (optimal) performance of the greedy algorithm degrades gracefully for valuations that are close to linear functions, provided these valuations are also approximately *submodular* (in a stronger than pointwise sense). (As noted above, some assumption beyond near-linearity is necessary for any positive result.) The standard arguments for proving approximation bounds for the greedy algorithm (e.g. [26]) do not imply this result, and we develop a new analysis for this purpose.

We then consider welfare maximization with demand oracles,[5] and find that welfare guarantees tend to be more robust in this model. First, the *value* of optimal social welfare can be approximated using demand queries within a factor of $\gamma_{\mathcal{C}} + \epsilon$, whenever the valuations are pointwise ϵ-close to a class $\mathcal{C}$ such that the integrality gap of the "configuration LP" is $\gamma_{\mathcal{C}}$. Since the configuration LP is the primary vehicle for developing approximation algorithms in the demand oracle model, we recover pointwise robustness essentially for all known approximation results in the demand oracle model (in terms of the optimal value). Another question, though, is whether an allocation achieving good welfare can be found in a computationally-efficient way. For some classes of valuations, we show that this is possible (and thus we achieve a $(1-\epsilon)$-approximation for valuations ϵ-close to linear, and a $(1-1/e-\epsilon)$-approximation for valuations ϵ-close to XOS). We remark that no extra assumption of near-submodularity is required here; this highlights another difference between the value and demand oracle models.

Another approach to finding an optimal allocation assuming the gross substitutes property is the Kelso-Crawford algorithm, also known as the *tâtonnement* procedure. In general, this procedure requires demand queries and thus also falls within the umbrella of the demand oracle model. Here we show that the performance of the Kelso-Crawford algorithm degrades smoothly for certain classes of functions, namely for valuations close to linear, close to unit-demand with $\{0,1\}$ values, and more generally close to transversal valuations (rank functions of a partition matroid). Unfortunately the Kelso-Crawford algorithm is not going to be a universally robust solution for general gross substitutes, either. As we show, its performance degrades discontinuously for valuations close to unit-demand (with unrestricted values), a seemingly minor extension of the cases above where we showed positive results.

In addition we show a counterexample to an approach based on Murota's cycle canceling algorithm, the remaining known way to solve the welfare maximization problem under the gross substitutes property.

[5] As noted above, with valuations that only satisfy gross substitutes approximately, demand oracles are substantially more powerful than value oracles.

In summary, the idea that approximation guarantees for various classes of valuations should degrade gracefully under small deviations from the class should be viewed with skepticism. Our negative results do not imply that that the folklore belief that “close to substitutes is easy” is wrong, but they do imply that there is no “generic reason” for why this might be the case. Our positive results show that robust guarantees can still be obtained in certain cases, but typically this requires additional care and often new ideas on top of known techniques.

1.2 Related Work and Organization

There is growing interest in the algorithmic game theory literature in the class of gross substitutes valuations prominent in the economic literature [36, 21, 37]. One reason for this is that welfare maximization for gross substitutes has deep mathematical and algorithmic roots – the simple subcase of unit-demand valuations already subsumes bipartite matching [39]. The main algorithmic techniques include ascending-price auctions [24], combinatorial cycle-canceling algorithms that utilize discrete convexity [31, 32], and linear programming [6, 33]. Submodular valuations are a strict superclass of gross substitutes, and for these the following is known: exact welfare maximization is hard, a 2-approximation can be achieved greedily [26], and continuous greedy – the optimal approximation algorithm in the value oracle model – achieves a $(1 - 1/e)$-approximation [46, 25]. In the demand query model, a 2-approximation is achieved via an ascending-price auction [17], and the $1 - 1/e$ barrier can be broken [15].

Not much previous work studies the extent to which welfare maximization is robust to small deviations. The closest result to ours is a theorem of [19], ruling out a constant-factor approximation when maximizing valuations that are pointwise ϵ-close to submodular, rather than to gross substitutes. Our Proposition 4 strengthens this negative result to apply to gross substitutes valuations and in fact even to valuations that are close to linear. Related lower bounds appear e.g. in [30, 38, 20], often relying on the same approach of hiding structure in a seemingly “nice” set function. The challenge in our case is to allow the function to appear linear.

Several related works differ from ours in the model of deviation, valuation classes considered, and/or setting: [19] studies welfare maximization with submodular valuations and random rather than adversarial noise. [29] introduces a distance metric from matroid rank functions in a single-parameter rather than combinatorial setting. [13] and references within study deviations from submodularity captured by graphical representations. [27] tests the convergence and performance of a heuristic approach based on cycle-canceling for subclasses of submodular valuations.

Finally, it is interesting to compare our work to [23], which studies divisible items and budgets rather than indivisible items and quasi-linear utilities. [23] defines an *approximate weak gross substitutes* property and shows that with this property a known auction mechanism guarantees each bidder approximately his demand (in a different sense than our Definition 22(4)). The conclusion of [23], by which “markets do not suddenly become intractable if they slightly violate the weak gross substitutes property”, is quite different from ours, showing an intriguing discrepancy between the divisible and indivisible models.

Organization. After preliminaries in Section 2, Section 3 presents two cautionary tales as to what can go wrong when simple valuations undergo smalls perturbations. We then establish positive results in the value oracle model (Section 4) and in the demand oracle model (Section 5), and conclude in Section 6. We defer some proofs to Appendices A–C. Negative results for standard algorithms appear in Appendix D and in the full version [40].

2 Preliminaries

A *combinatorial auction* (or *market*) includes a set of n players (or *bidders*) N and a set of m indivisible items M. We call a subset $S \subseteq M$ of items a *bundle*. For an ordered set of items $S = (s_1, s_2, \ldots, s_k)$ and for $j \in [k]$, we use the notation S_j to denote the "prefix" subset $\{s_1, \ldots s_{j-1}\}$ of items that appear before s_j (in particular, $S_1 = \emptyset$).

Every player i has a *valuation* $v_i : 2^M \to \mathbb{R}_+$ over bundles. Valuations are assumed to be monotone unless stated otherwise, i.e., $v(S) \leq v(T)$ for every two bundles $S \subseteq T$. Valuations are not necessarily normalized, i.e., $v(\emptyset)$ is not always equal to zero (normalization is not without loss of generality in our model of perturbation). [6] For every two bundles S, T we denote by $v(S \mid T)$ the *marginal* value $V(S \cup T) - v(T)$ of S given T.

An *allocation* $\mathcal{S}$ is a partition of the items in M into n bundles $S_1, \ldots, S_n$ where S_i is the allocation of player i. In a *full* allocation every item is allocated. The social efficiency of $\mathcal{S}$ is measured by its *welfare* $W(\mathcal{S}) = \sum_{i=1}^n v_i(S_i)$.

A *price vector* $p \in \mathbb{R}_+^m$ is a vector of item prices. We denote by $p(S)$ the aggregate price $\sum_{j \in S} p(j)$ of the bundle S. Given p, each player wishes to maximize his *quasi-linear utility*, i.e., to receive a bundle S maximizing $v_i(S) - p(S)$, which we say is *in demand*. If the player is already allocated a bundle T, he wishes to add a bundle $S \subseteq (M \setminus T)$ maximizing $v_i(S \mid T) - p(S)$, which we say is in demand *given* T. A bundle S is *individually rational (IR)* for a player i if $v_i(S) \geq p(S)$; it is *strongly IR* if every subset $T \subseteq S$ is IR for i.

A full allocation $\mathcal{S}$ and price vector p form a *Walrasian equilibrium* if for every player i, S_i is in i's demand given p. By the *1st welfare theorem*, $\mathcal{S}$ maximizes welfare.

Since valuations are in general of exponential size in m, the standard assumption is that they are accessed via *value oracles*, which return the value of any bundle upon query. In Sections 3.2 and 5 we discuss *demand oracles*, which given a price vector p, return a bundle in the player's demand given p. We allow p to include negative prices so that a demand oracle can return a bundle in the player's demand *given a previous allocation* T.

2.1 Valuation Classes

A valuation ℓ is *linear* (also known as affine) if there exists a vector $(l_1, \ldots, l_m) \geq 0$ and a scalar $c \geq 0$ such that $\ell(S) = c + \sum_{j \in S} l_j$ for every bundle S; it is *additive* if $c = 0$. A valuation r is *unit-demand* if there exists a vector $(\rho_1, \ldots, \rho_m) \geq 0$ such that $r(S) = \max_{j \in S} \rho_j$ for every bundle S.

Let $\mathcal{M} = (M, \mathcal{I})$ be a matroid over the ground set of items M, where $\mathcal{I}$ is the family of *feasible* bundles. [7] A maximal feasible set is called a *basis*, and the matroid *rank function* maps bundles to the cardinality of their largest feasible subset. Given a vector $(w_1, \ldots, w_m) \geq 0$ of item weights, the corresponding *weighted* rank function maps bundles to the weight of their heaviest feasible subset. A valuation r is an unweighted (resp., weighted)

[6] A natural example of a non-normalized valuation is that of a firm over sets of workers, where some positions are already filled [34].

[7] A set system $\mathcal{M} = (M, \mathcal{I})$ is a matroid if the following properties hold: (i) $\mathcal{I}$ is non-empty; (ii) $\mathcal{I}$ is downward-closed, that is, if $S \subseteq T$ and $T \in \mathcal{I}$ then $S \in \mathcal{I}$; (iii) for every $S, T \in \mathcal{I}$ such that $|S| < |T|$, there exists $t \in T \setminus S$ such that $S \cup \{t\} \in \mathcal{I}$ (see, e.g., [35]).

matroid rank function if there exists a matroid $\mathcal{M}$ such that r is its (weighted) rank function. Both additive and unit-demand valuations are types of weighted matroid rank functions.

A valuation v is *gross substitutes* if for every two price vectors $p \leq q$, for every bundle S in demand given p, there exists a bundle T in demand given q, which contains every item $j \in S$ for which $q(j) = p(j)$. [8] A valuation v is *submodular* if for every two bundles $S \subseteq T$ and item $j \notin T$, $v(j \mid S) \geq v(j \mid T)$, i.e., the marginal value of j is decreasing. A valuation v is *XOS* (also known as *fractionally subadditive*) if there exist additive valuations $a_1, a_2, \ldots$ such that $v(S) = \max_k a_k(S)$ for every bundle S. It is well-known that gross substitutes $\subset$ submodular $\subset$ XOS.

2.2 Welfare Maximization

There are three main algorithmic approaches to finding a welfare-maximizing allocation in combinatorial auctions with gross substitutes: (1) auction-based, (2) LP-based, and (3) combinatorial. For linear valuations there is also (4) greedy. The first two approaches use demand queries and run in polynomial time, while the latter two use value queries and run in strongly-polynomial time. We describe here Approaches (1) and (4); see also Appendix A.

In the Kelso-Crawford auction (Approach (1), Algorithm 1 in Appendix A), in each round an arbitrary player adds to his existing allocation a bundle in his demand given his current bundle and the current prices (with a δ-increase in the prices of items not currently in his allocation). The algorithm terminates when no player wants to add to his allocation. By the definition of gross substitutes, at any point in the algorithm every player's allocation is a subset of a bundle in his demand, and so:

▶ **Proposition 1** ([24]). *For gross substitutes valuations, Algorithm 1 converges to a (welfare-maximizing) Walrasian equilibrium as $\delta \to 0$.*

The greedy algorithm (Approach (4), Algorithm 2 in Appendix A) finds a value-maximizing bundle subject to a feasibility constraint $\mathcal{I}$ on the bundles;[9] this is a problem to which welfare maximization is well-known to reduce (see proof of Corollary 16). The greedy algorithm starts with an empty bundle, and in each iteration adds the item with maximum marginal value among all items that maintain feasibility, breaking ties arbitrarily.

▶ **Proposition 2** ([16]). *If v is linear and $\mathcal{I}$ is such that $\mathcal{M} = (M, \mathcal{I})$ is a matroid, then Algorithm 2 is optimal.*

3 Two Cautionary Tales

Do the nice properties of a valuation class degrade gracefully as one moves outside the class? This section describes two cautionary tales demonstrating that the answer can subtly depend on the notion of "being close," on the tractable class under consideration, and on the model of access to the valuations. Arguably the most natural general notion of "closeness" is pointwise:[10]

[8] The intuition behind this definition is revealed by the Kelso-Crawford algorithm below.
[9] The constraint $\mathcal{I}$ is simply a family of feasible bundles; we assume it is given by a feasibility oracle.
[10] We use relative error in the interest of scale-invariance (the "units" in which valuations are specified do not matter). However it is interesting to note that our negative results in this section would hold even for additive (rather than linear) valuations in a model of additive (rather than multiplicative) error, as studied e.g. by [10].

▶ **Definition 3.** A valuation $\tilde{v}$ is *ϵ-close* to submodular (or linear, gross substitutes, etc.), if there is a submodular (or linear, gross substitutes, etc.) valuation v such that $v(S) \leq \tilde{v}(S) \leq (1+\epsilon)v(S)$ for all bundles S.

3.1 Close-to-additive vs. Close-to-linear Valuations

We next show that approximate welfare maximization is hard for valuations that are ϵ-close to linear (a restrictive subclass of gross substitutes valuations).

▶ **Proposition 4.** *Given value oracles for valuations ϵ-close to linear, no algorithm using a subexponential number of queries can approximate the value of optimal social welfare within a factor better than polynomial in m, n.*

Proof Sketch. Let $|M| = m = an$ and let $(A_1, A_2, \ldots, A_n)$ be a random partition of M such that $|A_i| = a$. We define linear valuations as follows: $\ell_i(S) = \epsilon + \frac{1}{a}|S \cap A_i|$. We also define $w_i(S) = (1+\epsilon)\epsilon + \frac{1}{an}|S|$. Note that if we do not know the partition $(A_1, \ldots, A_n)$, which is randomized, $\ell_i(S)$ will be close to its expectation which is $w_i(S)$. Let us define the following valuations v_i:

- If $\big||S \cap A_i| - \frac{1}{n}|S|\big| \leq \epsilon^2 a$, then $v_i(S) = w_i(S)$.
- If $\big||S \cap A_i| - \frac{1}{n}|S|\big| > \epsilon^2 a$, then $v_i(S) = \ell_i(S)$.

Note that in the first case, we have $v_i(S) \geq (1+\epsilon)\epsilon + \frac{1}{a}(|S \cap A_i| - \epsilon^2 a) = \ell_i(S)$ and $v_i(S) \leq (1+\epsilon)\epsilon + \frac{1}{a}(|S \cap A_i| + \epsilon^2 a) \leq (1+2\epsilon)\epsilon + \frac{1}{a}|S \cap A_i| \leq (1+2\epsilon)\ell_i(S)$. So v_i is 2ϵ-close to ℓ_i.

By Chernoff bounds, for a fixed query S, the probability that $\big||S \cap A_i| - \frac{1}{n}|S|\big| > \epsilon^2 a$ is $e^{-\Omega(\epsilon^4 a)}$. Hence, with high probability we are always in the first case above, the returned value depends only on $|S|$ and hence the algorithm does not learn any information about A_i. Therefore (by standard arguments), it would require exponentially many queries to find any set such that $\big||S \cap A_i| - \frac{1}{n}|S|\big| > \epsilon a$ and hence distinguish whether the input valuations are v_i or w_i.

The optimal solution under v_i is $S_i = A_i$ which gives welfare $\sum_{i=1}^n v_i(A_i) > \frac{1}{a}\sum_{i=1}^n |A_i| = n$, while the optimal welfare under w_i is $(1+\epsilon)\epsilon n + 1$. So the approximation factor cannot be better than $(1+\epsilon)\epsilon + \frac{1}{n}$.

Note that we can set $\epsilon = 1/a^{1/4-\delta} = (n/m)^{1/4-\delta}$ and the high probability statements still hold. Therefore, we can push the hardness factor to $\max\{(n/m)^{1/4-\delta}, 1/n\}$. ◀

This hardness result relies strongly on the value oracle model – things are quite different in the demand oracle model, as we discuss in Section 3.2 below.

Proposition 4 is a startling contrast to the case of valuations that are ϵ-close to *additive* valuations, which differ only by requiring the empty set to have value 0. Here, approximate welfare maximization is easy (a proof appears for completeness in Appendix B).

▶ **Proposition 5.** *Welfare maximization can be solved within a factor of $1+\epsilon$ for ϵ-close to additive valuations.*

The proof of Proposition 5 makes the more general point that, whenever the approximating "nice" valuation $\tilde{v}$ can be (approximately) recovered from the given valuation v, then welfare approximation guarantees carry over (applying an off-the-shelf approximation algorithm to the valuations $\tilde{v}$). As Proposition 4 makes clear, however, in many cases it is not possible to efficiently reconstruct an approximating "nice" valuation, even under the promise that such a valuation exists.

3.2 Value Queries vs. Demand Queries

One remarkable property of gross substitutes valuations is that there is no difference between the value oracle and demand oracle models: Demand queries can be simulated efficiently by value queries (via the greedy algorithm), and hence any algorithm in the demand oracle model can also be implemented in the value oracle model ([36] and references within). Unfortunately, this is no longer true for valuations that are ϵ-close to gross substitutes, or even ϵ-close to linear.

► **Proposition 6.** *Given a value oracle to a valuation ϵ-close to linear, answering demand queries requires an exponential number of value queries.*

This can be proved by a construction similar to the one above. However, let us instead present an indirect argument which shows more.

► **Lemma 7.** *For any class of valuations $\mathcal{C}$ such that the integrality gap of the configuration LP is $\gamma \geq 1$, it is possible to estimate the optimal social welfare within a multiplicative factor of $(1+\epsilon)\gamma$ for valuations that are ϵ-close to $\mathcal{C}$ in the demand oracle model.*

Proof. Consider the configuration LP:

$$\begin{aligned} \max \quad & \textstyle\sum_{i=1}^{n} \sum_{S \subseteq [m]} v_i(S) x_{i,S} : \\ \forall j \in [m]; \quad & \textstyle\sum_{i=1}^{n} \sum_{S : j \in S} x_{i,S} \leq 1, \\ \forall i \in [n]; \quad & \textstyle\sum_{S \subseteq [m]} x_{i,S} = 1, \\ & x_{i,S} \geq 0. \end{aligned}$$

Let us denote by $\mathrm{LP}(\mathbf{v})$ and $\mathrm{OPT}(\mathbf{v})$ the LP optimum and optimal social welfare, respectively, under valuations $\mathbf{v}$. The assumption is that for valuations $\mathbf{v} \in \mathcal{C}$, $\mathrm{OPT}(\mathbf{v}) \leq \mathrm{LP}(\mathbf{v}) \leq \gamma \cdot \mathrm{OPT}(\mathbf{v})$. Let us solve the LP for valuations $\tilde{\mathbf{v}}$ that are pointwise ϵ-close to valuations $\mathbf{v} \in \mathcal{C}$. (This is possible since we assume that we have demand oracles for $\tilde{\mathbf{v}}$, which gives a separation oracle for the dual; see [33].) We have $v_i(S) \leq \tilde{v}_i(S) \leq (1+\epsilon) v_i(S)$ for each i and S. Therefore, the same inequalities hold for the LP optimum as well as optimal social welfare, and we obtain $\mathrm{OPT}(\tilde{\mathbf{v}}) \leq \mathrm{LP}(\tilde{\mathbf{v}}) \leq (1+\epsilon)\,\mathrm{LP}(\mathbf{v}) \leq (1+\epsilon)\gamma \cdot \mathrm{OPT}(\mathbf{v}) \leq (1+\epsilon)\gamma \cdot \mathrm{OPT}(\tilde{\mathbf{v}})$. ◄

Since the configuration LP is known to be integral for gross substitutes valuations (see e.g. [45]), we obtain the following.

► **Corollary 8.** *For valuations ϵ-close to gross substitutes, it is possible to estimate the optimal social welfare to within a multiplicative factor of $1+\epsilon$ in the demand oracle model.*

This implies Proposition 6: a simulation of demand queries would allow us to approximate social welfare within $1+\epsilon$ for valuations ϵ-close to gross substitutes and as a special case ϵ-close to linear. We know from Proposition 4 that this (and even much weaker approximations) would require exponentially many value queries. Actually we obtain a stronger statement: It is not possible to answer demand queries via value queries for ϵ-close to linear valuations even approximately, under *any notion of approximation* that would be useful for approximating the optimal social welfare. This is because, again, such a simulation would lead to a $(1+\epsilon)$-approximation of social welfare via value queries, which Proposition 4 rules out.

4 Positive Results for Restricted Valuation Classes

Motivated by the hardness results in the previous section for valuations ϵ-close to linear in the value oracle model, this section considers stronger notions of closeness. In Section 4.1 we define a semantic notion of *marginal* closeness. In Section 4.2 we analyze the greedy algorithm for ϵ-close to linear valuations, but this time to avoid the impossibility results we assume they are also marginally close to submodular. Intuitively, such valuations enable positive results by approximately maintaining some of the semantic properties of linear valuations (namely submodularity).

4.1 Marginal Closeness

Recall that a pointwise approximation of a valuation is a set function with approximately the same output as the valuation for every input. A natural strengthening of this is to have the set function's discrete derivatives approximate those of the valuation – i.e., the items' marginal values. For every item j, its marginal value $v(j \mid \cdot)$ is a set function mapping S to $v(j \mid S)$, and together these set functions encode important properties of the valuation. For example, linear valuations can be characterized as valuations for which the marginal of every item is a constant function; submodular valuations can be characterized as valuations for which the marginal of every item is a non-increasing function; and unweighted matroid rank functions can be characterized as valuations for which the marginal of every item is a non-increasing zero-one function [41].

We remark that even without any strengthening, pointwise ϵ-closeness has the following useful implication regarding the "closeness" of the marginal values of bundles:

▶ **Observation 9.** *For every valuation v that is ϵ-close to a valuation v', and for every two bundles $S, T \subseteq M$, $v'(S \mid T) - \epsilon v'(T) \leq v(S \mid T) \leq v'(S \mid T) + \epsilon v'(S \cup T)$.*

Proof. By the definition of ϵ-closeness, $v(S \mid T) = v(S \cup T) - v(T) \geq v'(S \cup T) - (1+\epsilon)v'(T) = v'(S \mid T) - \epsilon v'(T)$. The upper bound follows similarly. ◀

Observation 9 will be useful in proving some of our positive results (see Sections 4.2 and 5.2). Yet its guarantee is relatively weak: the "error terms" depend on $v'(T)$ or $v'(S \cup T)$, and so can be large if these terms are large. We now define the stronger notion of marginal pointwise approximation:[11]

▶ **Definition 10.** A valuation v is *marginal-ϵ-close* to a class $\mathcal{C}$ of marginals (constant, non-increasing, etc.), if for every item j the corresponding marginal value function $v(j \mid \cdot)$ is (pointwise) ϵ-close to a function $g_j \in \mathcal{C}$.

Let us now see what semantic properties follow from this definition. Consider the class of valuations that are marginal-ϵ-close to decreasing. It turns out that this coincides with a previously-studied class of approximately submodular valuations (see Appendix B for missing proofs).

▶ **Definition 11** ([26]). Let $\alpha \geq 1$. A valuation v is *α-submodular* if for every two bundles $S \subseteq T$ and item $j \notin T$, $\alpha v(j \mid S) \geq v(j \mid T)$.

▶ **Proposition 12.** *A valuation v is marginal-ϵ-close to the class of decreasing functions if and only if it is $(1+\epsilon)$-submodular.*

[11] This notion is equivalent to having erroneous oracle access to marginal values, an idea mentioned in [9].

Since we are interested in closeness notions that maintain semantic properties, in the next section we consider what happens when the ϵ-close to linear valuations from Section 3 are also α-submodular. Such valuations are well-studied in the literature as they contain submodular valuations with bounded curvature, which were introduced by [11] and arise frequently in the optimization and learning literatures ([42, 44] and references within):

▶ **Definition 13** **([11, 44]).** A valuation v has *curvature* $c \in [0,1]$ if for every item j and bundle S such that $j \notin S$, $v(j \mid S) \geq (1-c)v(j \mid \emptyset)$.

▶ **Proposition 14.** *An α-submodular valuation v with curvature c is $\frac{\alpha-1+c}{1-c}$-close to a linear valuation.*

4.2 The Greedy Algorithm

Recall that the greedy algorithm (Algorithm 2) is optimal for linear valuations; in this section and in Appendix B we prove the following robustness theorem, whose implication for welfare maximization is stated in Corollary 16.

▶ **Theorem 15.** *Let v be an α-submodular valuation that is ϵ-close to a linear valuation ℓ. Let $\mathcal{M} = (M, \mathcal{I})$ be a matroid of rank k (represented by an independence oracle). The greedy algorithm returns an independent set $S \in \mathcal{I}$ such that $v(S) \geq \frac{1-3\epsilon}{\alpha} v(S^*)$, where $S^* \in \arg\max_{I \subseteq \mathcal{I}} v(I)$.*

▶ **Corollary 16.** *In a market with α-submodular valuations $v_1, \ldots, v_n$ that are ϵ-close to linear valuations $\ell_1, \ldots, \ell_n$, there is a polynomial time algorithm with value access that finds an allocation with $\frac{1-3\epsilon}{\alpha}$-approximately optimal welfare.*

Proof. As in [26], we define a valuation v over player-item pairs, such that for a set S of such pairs, $v(S) = \sum_i v_i(S_i)$ where S_i is the set of items paired with player i in S. We show that v is also α-submodular and ϵ-close to linear: It is ϵ-close to linear since $\sum_i \ell_i(S_i) \leq v(S) \leq (1+\epsilon)\sum_i \ell_i(S_i)$, and the sum of linear valuations is linear. It is α-submodular since for every $S \subseteq T$ and $(i,j) \notin T$, $\alpha v((i,j) \mid S) = \alpha v_i(j \mid S_i) \geq v_i(j \mid T_i) = v((i,j), T)$. The proof follows by applying Theorem 15 to v and to the partition matroid $\mathcal{M}$ over player-item pairs, which allows each item to be paired with no more than one player. ◀

The following lemma (whose proof appears in Appendix B) relates the sum of marginals to the linear contribution, and is applied in the proof of Theorem 15. Recall that v is ϵ-close to the linear valuation ℓ, and that for an ordered set of items $X = (x_1, x_2, \ldots, x_k)$, X_j denotes the prefix $\{x_1, \ldots x_{j-1}\}$.

▶ **Lemma 17.** *Let X and Z be two disjoint sets. X is ordered and has m_X items. Let $Y_1, \ldots, Y_{m_X}$ be sets such that $Y_j \subseteq X_j \cup Z$ for every j. Then $\alpha \sum_{j=1}^{m_X} v(x_j \mid Y_j) \geq \sum_{j=1}^{m_X} \ell(x_j) - \epsilon\ell(Z)$.*

Before proving Theorem 15, we highlight the difference between our proof and the standard item-by-item analysis of the greedy algorithm (see, e.g., [26]). The standard analysis shows that greedily choosing the next item maintains the optimal value up to a small error. However this does not provide a good approximation factor in our case because of the distortion caused by the pointwise errors. For example, greedy may choose to include an item j from the optimal solution, but at the particular stage in which j is included its marginal value may not be high relative to its linear contribution (due to the error term in the marginal value shown in Observation 9). This suggests analyzing multiple items at a time so that the error terms will average out. We split the items chosen by greedy into two sets – those that coincide with the optimal solution and the rest – and analyze each set as a whole in order to establish the claimed approximation ratio.

Proof of Theorem 15. Denote by S the independent set that greedy returns given the valuation v and matroid $\mathcal{M}$, ordered by the order in which the items were greedily added. Denote by S^* an independent set with optimal value $v(S^*)$ subject to the matroid constraint. Due to monotonicity we can assume without loss of generality that the sizes of both S and S^* are k, i.e., that both sets are bases.

Let $P = \{j \mid s_j \in S \cap S^*\}$ be the positions in S of the items that also appear in S^*, and let Q be the remaining positions in S. Order S^* such that the items in $S \cap S^*$ are in positions P. The rest of the items in S^* are ordered among the remaining positions in S^* in the following way: By the exchange property of matroids, there exists a bijection f from $S \setminus S^*$ to $S^* \setminus S$, such that for every position $j \in Q$, $S \setminus \{s_j\} \cup \{f(s_j)\}$ is independent [41]. For every $j \in Q$ we set $s^*_j = f(s_j)$. This ensures that s^*_j can be added to S_j without violating feasibility. We use the notation $S(P)$ and $S(Q)$ for the set of items in S in positions P and Q, respectively; similarly for $S^*(P)$ and $S^*(Q)$.

By the definition of greedy, and because adding s^*_j to S_j maintains feasibility for every $j \in Q$, it holds that $v(s_j \mid S_j) \geq v(s^*_j \mid S_j)$ for every $j \in [k]$. Thus

$$v(S) \geq v(\emptyset) + \sum_{j \in P} v(s_j \mid S_j) + \sum_{j \in Q} v(s^*_j \mid S_j). \tag{1}$$

We now invoke Lemma 17 twice. By Lemma 17 instantiated with $X = S^*(Q)$ and $Z = S$ (note these are indeed disjoint), and using that $S_j \subseteq S$, we get that

$$\alpha \sum_{j \in Q} v(s^*_j \mid S_j) \geq \sum_{j \in Q} \ell(s^*_j) - \epsilon \ell(S). \tag{2}$$

By Lemma 17 instantiated with $X = S(P)$ and $Z = S(Q)$ (these are also disjoint), and using that $S_j \subseteq S(Q) \cup \{s_{p_1}, \ldots, s_{p_{j-1}}\}$, we get that

$$\alpha \sum_{j \in P} v(s_j \mid S_j) \geq \sum_{j \in P} \ell(s_j) - \epsilon \ell(S) = \sum_{j \in P} \ell(s^*_j) - \epsilon \ell(S). \tag{3}$$

Combining (1), (2) and (3), we get that $\alpha v(S) \geq \alpha v(\emptyset) + \sum_{j=1}^{k} \ell(s^*_j) - 2\epsilon \ell(S)$. Since v is ϵ-close to ℓ and S^* is optimal for v, $\ell(S) \leq v(S) \leq v(S^*)$. Again using v's closeness to ℓ, $\alpha v(\emptyset) \geq \alpha c \geq c$. Putting these together we get that $\alpha v(S) \geq \ell(S^*) - 2\epsilon v(S^*) \geq v(S^*)/(1+\epsilon) - 2\epsilon v(S^*) \geq (1 - 3\epsilon) v(S^*)$, as required. ◀

5 The Demand Oracle Model

In this section we achieve positive results by considering a stronger oracle model.

5.1 Rounding the Configuration LP

By Lemma 7 and Corollary 8, we already know that polynomial-time *estimation* of the optimal welfare is robust in the demand oracle model. A different question is whether *finding* an allocation given a fractional solution of the configuration LP is also robust. Here we observe that in some cases, existing rounding techniques yield the result that we want.

The following proposition does not assume submodularity:

▶ **Proposition 18.** *For combinatorial auctions in the demand oracle model, there is a polynomial time algorithm that finds:*

- *A $(1-\epsilon)$-approximately optimal allocation, if the valuations are ϵ-close to linear;*
- *A $(1-1/e-\epsilon)$-approximately optimal allocation, if the valuations are ϵ-close to XOS (or in particular ϵ-close to submodular).*

Proof. We solve the configuration LP using demand queries. Let $\tilde{v}_i$ denote the true valuations and v_i the linear/XOS valuations that the $\tilde{v}_i$ are close to.

In the case of valuations close to linear, we round the fractional solution by allocating each item j to player i with probability $y_{ij} = \sum_{S:j\in S} x_{i,S}$. Call the resulting random sets R_i. For the underlying linear valuations v_i, it is clearly the case that $\mathbb{E}[v_i(R_i)] = \sum_S x_{i,S} v_i(S)$. Therefore, $\sum_i \mathbb{E}[\tilde{v}_i(R_i)] \geq \sum_i \mathbb{E}[v_i(R_i)] = \sum_{i,S} x_{i,S} v_i(S) \geq \frac{1}{1+\epsilon}\,\mathrm{LP}(\tilde{v}_i) \geq (1-\epsilon)\,\mathrm{OPT}$.

For valuations close to XOS, we allocate a set S tentatively to player i with probability $x_{i,S}$, and then we use the contention resolution technique to resolve conflicts [12, 15]. This technique has the property that conditioned on requesting an item j, a player receives it with conditional probability at least $1-1/e$. Denote by R_i the random set that player i receives after contention resolution. Using the fractional subadditivity of XOS functions, we obtain that $\mathbb{E}[v_i(R_i)] \geq (1-1/e)\sum_S x_{i,S} v_i(S)$. Hence, similarly to the case above, $\sum_i \mathbb{E}[\tilde{v}_i(R_i)] \geq \sum_i \mathbb{E}[v_i(R_i)] \geq (1-1/e)\sum_{i,S} x_{i,S} v_i(S) \geq \frac{1-1/e}{1+\epsilon}\,\mathrm{LP}(\tilde{v}_i) \geq (1-1/e-\epsilon)\,\mathrm{OPT}$. ◀

5.2 Positive Results for Kelso-Crawford

The Kelso-Crawford algorithm (Algorithm 1) uses demand queries in order to find, in each iteration, the items to add to a player's existing bundle. Here we show that it can work well for markets in which the valuations are approximately submodular (maintain semantic closeness) and ϵ-close to simple subclasses of gross substitutes. Our high-level approach is to show that Kelso-Crawford finds an allocation which achieves approximately optimal welfare, as well as a price vector that is "approximately stabilizing" (in a sense made precise in Definition 22, which may be of independent interest).

▶ Remark. For simplicity, our analysis of the Kelso-Crawford algorithm shall treat prices as if raised continuously rather than discretely. This means that the approximation factors we state in this section hold up to a small additive error.[12]

We first establish that the Kelso-Crawford algorithm works well for the class of α-submodular and ϵ-close to linear valuations, as studied above in Section 4.

▶ **Theorem 19** (Linear). *In a market with α-submodular valuations that are ϵ-close to linear valuations, the Kelso-Crawford algorithm finds an allocation with $\frac{1}{\alpha+2\epsilon}$-approximately optimal welfare (up to a vanishing error).*

On the negative side and somewhat surprisingly, we find this positive result does not extend even to ϵ-close to unit-demand (and still α-submodular) valuations (see Appendix D and the full version):

▶ **Proposition 20.** *For every $\epsilon \leq 1$, there exists a market with $O(1/\epsilon)$ players whose submodular valuations are ϵ-close to unit-demand, for which the Kelso-Crawford algorithm with adversarial ordering of the players finds an allocation with $\approx 2/3$ of the optimal welfare.*

[12] If the prices are increased by discrete δ increments, then the additive error is of order $O(mn\delta)$.

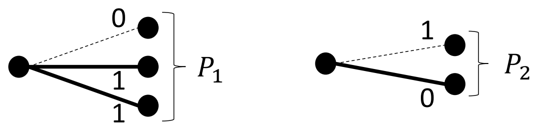

Figure 1 A graphical representation of an unweighted transversal valuation r. The edges correspond to items and have $\{0, 1\}$ weights; they are partitioned in this case into two parts P_1, P_2. The value that r attributes to the subset of solid bold edges in this example is 1, since this is the weight of the maximum-weight matching within the subset.

We thus proceed to consider the *unweighted* version of unit-demand valuations, and more generally the direct sums of such valuations – called *unweighted transversal* valuations. Such valuations arise in natural economic environments, e.g. in the context of labor markets. [13] For these valuations we establish in Theorem 21 that Kelso-Crawford maintains good welfare guarantees.

▶ **Theorem 21** (Transversal). *In a market with α-submodular valuations that are ϵ-close to unweighted transversal valuations, the Kelso-Crawford algorithm finds an allocation with $\frac{1}{\alpha(1+3\epsilon)^2}$-approximately optimal welfare (up to a vanishing error).*

In more detail, a valuation r is an unweighted unit-demand valuation if $v(S) = \max_{j \in S} r(j)$ for every bundle S and $r(j) \in \{0, 1\}$ for every item j. A valuation r is an unweighted transversal valuation if there exists a partition $\mathcal{P} = (P_1, \ldots, P_k)$ of the items such that $v(S) = \sum_{P \in \mathcal{P}} \max_{j \in S \cap P} r(j)$ for every bundle S, and $r(j) \in \{0, 1\}$ for every item j. An equivalent definition in terms of matchings in graphs is the following: r is an unweighted transversal valuation if it can be represented by a bipartite graph whose vertices on one side of the graph all have degree 1, and whose edges M have $\{0, 1\}$ edge weights. The value for a bundle of edges S is the weight of the maximum matching in the bipartite graph induced by S – see Figure 1 for an example.

5.2.1 Biased Walrasian Equilibrium: Definition and Properties

The proofs of Theorems 19 and 21 follow from the welfare guarantees of a solution concept that we call a "biased" Walrasian equilibrium (Definition 22 and Proposition 23), combined with an appropriate lemma showing that Kelso-Crawford converges to such an equilibrium for the relevant class of valuations (Lemma 27 for Theorem 19, and Lemma 29 for Theorem 21). Recall that the standard proof establishing the optimality of Kelso-Crawford for gross substitutes relies on showing convergence to a Walrasian equilibrium. Unfortunately in our case it does not hold – even approximately – that Kelso-Crawford allocates to each player a bundle in his demand. Instead, our proofs will define and utilize a different notion of an approximate Walrasian equilibrium.

▶ **Definition 22.** Consider a market with n players, and let $\mu \in [0, 1]$. A full allocation $\mathcal{S}$ and a price vector p form a *μ-biased Walrasian equilibrium* if there exists $\mu' \in [\mu, 1]$ such that for every alternative allocation T and for every player i,

$$\frac{\mu'}{\mu} v_i(S_i) - p(S_i) \geq \mu' v_i(T_i) - p(T_i). \tag{4}$$

[13] Consider for example a firm wishing to hire a team of several workers, each with a different specialization, from a pool of specialist workers who are each either acceptable to the firm or not. The firm's valuation over workers is unweighted transversal.

▶ **Proposition 23** (Approximate first welfare theorem.). *The welfare of a μ-biased Walrasian equilibrium is a μ-approximation to the optimal welfare.*

Proof. Let $\mathcal{S}^*$ be an optimal allocation. Without loss of generality we can assume that S^* is a full allocation. By summing up Inequality (4) over all players, we get $(\mu'/\mu)W(\mathcal{S}) - p(\mathcal{S}) \geq \mu' W(\mathcal{S}^*) - p(\mathcal{S}^*)$. Since both $\mathcal{S}, \mathcal{S}^*$ are full allocations, $p(\mathcal{S}) = p(\mathcal{S}^*)$. We conclude that $W(\mathcal{S}) \geq \mu \, \mathrm{OPT}$. ◀

A possible economic interpretation of a μ-biased Walrasian equilibrium is that its allocation and prices induce market stability under the *endowment effect* (see, e.g., [22]). This effect is modeled as an increase of factor μ'/μ in a player's value for the bundle he owns, and a decrease of factor μ' in his value for bundles he does not own.

5.2.2 Convergence to Biased Walrasian Equilibrium

To give intuition for establishing convergence to biased Walrasian equilibrium (see Lemmas 27 and 29), let us compare the case of linear valuations to that of ϵ-close to linear (and α-submodular) valuations. The analysis of Kelso-Crawford for the linear case is simple – every player ends up with the items for which he has the highest marginal value and can afford to pay the highest price. In the close-to-linear case, an item could end up belonging to the wrong player for two reasons: at some point, given his current allocation (which is subject to changes), either (1) the right player has a low marginal value for the item, or (2) the wrong player has a high marginal value for it. The latter issue is resolved by submodularity, but the former could drive Kelso-Crawford to performance that is bounded away from optimal, as demonstrated in Proposition 20. Thus the crux of the convergence proofs is to show this cannot happen for the classes of valuations in question. For the close-to-linear case (Lemma 27), we achieve this by considering the marginal value of all "under-valued" items together if we were to add them to the player's final allocation. The transversal case (Lemma 29) is more involved since we cannot analyze only the addition of items to the player's final allocation – we need to also consider swapping out items from the final allocation and replacing them with others. This case is deferred to Appendix C.

We conclude this subsection with a generalization of a well-known invariant property of Kelso-Crawford which holds for submodular valuations [17]. The following generalization to α-submodular valuations is used below to prove Lemmas 27 and 29.

▶ **Definition 24.** Let v be a valuation and p be a price vector. A bundle S is *α-IR* (α-individually rational) for v given p if $\alpha v(S) \geq p(S)$. A bundle S is *strongly α-IR* for v given p if $\alpha v(T) \geq p(T)$ for every $T \subseteq S$.

▶ **Lemma 25.** *Consider a market with a player whose valuation v is α-submodular. Then the Kelso-Crawford algorithm maintains the following invariant: the player's allocation throughout the algorithm is strongly α-IR.*

Proof. Since removing items from the player's allocation maintains the strong α-IR property, the only case we need to check is when a bundle T is added to the player's current allocation S. We know that T maximizes the player's utility given S and the current price vector p. We show this by induction. In the base case all prices are 0 and so the invariant holds. Now assume for contradiction that $S \cup T$ is not strongly α-IR, i.e., there exists a set $S' \cup T'$ where $S' \subseteq S$ and $T' \subseteq T$ such that $\alpha v(S' \cup T') < p(S') + p(T')$. By the induction assumption, $\alpha v(S') \geq p(S')$, and so it must be the case that $\alpha v(T' \mid S') < p(T')$. So using α-submodularity we can write the utility $v(T \mid S) - p(T)$ as $v(T \setminus T' \mid S) + v(T' \mid S \cup T \setminus T') - p(T \setminus T') - p(T') \leq v(T \setminus T' \mid$

$S) + \alpha v(T' \mid S') - p(T \setminus T') - p(T') < v(T \setminus T' \mid S) - p(T \setminus T')$. Thus adding $T \setminus T'$ to S adds more to the utility than adding T to S, contradiction. ◀

The following corollary of Lemma 25 generalizes a result of [17, Propositions 1-2].

▶ **Corollary 26.** *In a market with α-submodular valuations, the Kelso-Crawford algorithm finds an allocation with $(1+\alpha)$-approximately optimal welfare (up to a vanishing error).*

5.3 Kelso-Crawford and Close-to-Linear Valuations

▶ **Lemma 27.** *In a market with α-submodular valuations $v_1, \ldots, v_n$ that are ϵ-close to linear valuations $r_1, \ldots, r_n$, the Kelso-Crawford algorithm converges to a $\frac{1}{\alpha+2\epsilon}$-biased Walrasian equilibrium.*

Proof. We can assume without loss of generality that Kelso-Crawford returns a full allocation (see proof of Theorem 21). From now until the end of the proof, fix a player i. Let S_i be player i's allocation and let p be the price vector at termination of the KC algorithm. Let T_i be an alternative bundle for player i. We show that

$$v_i(S_i) - p(S_i) \geq v_i(T_i) - p(T_i) - (2\epsilon + \alpha - 1)\, v_i(S_i). \tag{5}$$

This is sufficient to complete the proof, since by summing up over all players and rearranging, we get

$$(2\epsilon + \alpha) \sum_i v_i(S_i) - \sum_i p(S_i) \geq \sum_i v_i(T_i) - \sum_i p(T_i),$$

and so $\mu \geq 1/(\alpha + 2\epsilon)$.

It remains to prove Inequality (5). For simplicity we omit i from the notation. Without loss of generality, assume that in the last round of Kelso-Crawford, the player added to his existing bundle, which we denote by B, a bundle which maximizes his utility given B and given the price vector p. So the player's utility at termination $v_i(S_i) - p(S_i)$ is at least $u_p(B \cup T)$. We use the notation $B' = B \setminus T$, $T' = T \setminus B$, and $C = B \cap T$. We can now write the lower bound on the player's utility as

$$u_p(B \cup T) = v(\emptyset) + v(B' \mid \emptyset) + v(C \mid B') + v(T' \mid B) - p(B') - p(C) - p(T'). \tag{6}$$

By Lemma 25, B is α-strongly IR. So $v(B' \mid \emptyset) + (\alpha - 1)v(B' \mid \emptyset) \geq p(B')$. By Observation 9, $v(C \mid B') \geq \sum_{j \in C} \ell(j) - \epsilon v(B)$. Again by Observation 9, $v(T' \mid B) \geq \sum_{j \in T'} \ell(j) - \epsilon v(B)$. Plugging in these inequalities to (6) and using v's ϵ-closeness to ℓ we get

$$\begin{aligned} u_p(B \cup T) &\geq c - (\alpha - 1)v(B' \mid \emptyset) + \sum_{j \in T} \ell(j) - p(T) - 2\epsilon v(B) \\ &\geq \ell(T) - p(T) - (2\epsilon + \alpha - 1)\, v(B). \end{aligned}$$

Inequality (5) follows, and this completes the proof. ◀

Proof of Theorem 19. The proof follows directly from Proposition 23 combined with Lemma 27. ◀

6 Conclusion

Let us summarize the findings of this paper: The robustness of results for various classes of valuations should not be assumed without further investigation. The default assumption should be that positive results are *not* robust and may break down abruptly when small deviations from the class in question are introduced. Robust results can be attained, but they are surprisingly challenging to obtain even in simple cases. Most importantly, it matters how "closeness" to a class is defined, and what the other attributes of the variant of the problem are (oracle model, class of valuations).

Acknowledgements. A preliminarily version appeared as [40].

References

1 Lawrence M. Ausubel, Peter Cramton, R. Preston McAfee, and John McMillan. Synergies in wireless telephony: Evidence from the broadband PCS auctions. *Journal of Economics and Management Strategy*, 6(3):497–527, 1997.
2 Lawrence M. Ausubel and Paul R. Milgrom. The lovely but lonely Vickrey auction. In Peter Cramton, Yoav Shoham, and Richard Steinberg, editors, *Combinatorial Auctions*, chapter 1, pages 57–95. MIT Press, Boston, MA, USA, 2006.
3 Eric Balkanski, Aviad Rubinstein, and Yaron Singer. The limitations of optimization from samples. Working paper, 2016.
4 Alexandre Belloni, Tengyuan Liang, Hariharan Narayanan, and Alexander Rakhlin. Escaping the local minima via simulated annealing: Optimization of approximately convex functions. In *Proceedings of the 28th Annual Conference on Learning Theory*, pages 240–265, 2015.
5 Dimitris Bertsimas and Aurélie Thiele. *Robust and Data-Driven Optimization: Modern Decision Making Under Uncertainty*, chapter 5, pages 95–122. INFORMS PubsOnline, 2014. TutORials in Operations Research.
6 Sushil Bikhchandani and John W. Mamer. Competitive equilibrium in an exchange economy with indivisibilities. *Journal of Economic Theory*, 74(2):385–413, 1997.
7 Liad Blumrosen and Noam Nisan. Combinatorial auctions. In Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani, editors, *Algorithmic Game Theory*, chapter 11. Cambridge University Press, 2007.
8 M. M. Bykowsky, R. J. Cull, and J. O. Ledyard. Mutually destructive bidding: The FCC auction design problem. *Journal of Regulatory Economics*, 17(3):205–228, 2000.
9 Gruia Calinescu, Chandra Chekuri, Martin Pál, and Jan Vondrák. Maximizing a monotone submodular function subject to a matroid constraint. *SIAM Journal on Computing*, 40(6):1740–1766, 2011.
10 Flavio Chierichetti, Abhimanyu Das, Anirban Dasgupta, and Ravi Kumar. Approximate modularity. In *Proceedings of the 56th Symposium on Foundations of Computer Science*, pages 1143–1162, 2015.
11 Michele Conforti and Gérard Cornuéjols. Submodular set functions, matroids and the greedy algorithm: Tight worst-case bounds and some generalizations of the Rado-Edmonds theorem. *Discrete Applied Mathematics*, 7(3):251–274, 1984.
12 Uriel Feige. On maximizing welfare when utility functions are subadditive. *SIAM J. Comput.*, 39(1):122–142, 2009.
13 Uriel Feige, Michal Feldman, Nicole Immorlica, Rani Izsak, Brendan Lucier, and Vasilis Syrgkanis. A unifying hierarchy of valuations with complements and substitutes. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, pages 872–878, 2014.
14 Uriel Feige, Michal Feldman, and Inbal Talgam-Cohen. Approximate modularity revisited. In *Proc. of the 49th Annual ACM Symp. on Theory of Computing*, 2017. To appear.

15 Uriel Feige and Jan Vondrák. The submodular welfare problem with demand queries. *Theory of Computing*, 6(1):247–290, 2010.

16 M. L. Fisher, G. L. Nemhauser, and L. A. Wolsey. An analysis of approximations for maximizing submodular set functions – II. *Mathematical Programming Study*, 8:73–87, 1978.

17 Hu Fu, Robert Kleinberg, and Ron Lavi. Conditional equilibrium outcomes via ascending price processes with applications to combinatorial auctions with item bidding. In *Proceedings of the 13th ACM Conference on Economics and Computation*, page 586, 2012. Extended abstract.

18 Faruk Gul and Ennio Stacchetti. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory*, 87:95–124, 1999.

19 Avinatan Hassidim and Yaron Singer. Submodular optimization under noise. Manuscript, 2016.

20 John William Hatfield, Nicole Immorlica, and Scott Duke Kominers. Testing substitutability. *Games and Economic Behavior*, 75(2):639–645, 2012.

21 Justin Hsu, Jamie Morgenstern, Ryan Rogers, Aaron Roth, and Rakesh Vohra. Do prices coordinate markets? To appear in STOC 2016, 2016.

22 Daniel Kahneman, Jack L. Knetsch, and Richard H. Thaler. Experimental tests of the endowment effect and the Coase theorem. *Journal of Political Economy*, 98(6):1325–1348, 1990.

23 Chinmay Karande and Nikhil R. Devanur. Computing market equilibrium: Beyond weak gross substitutes. In *Proceedings of the 3rd International Workshop on Internet and Network Economics*, pages 368–373, 2007.

24 A. Kelso and V. Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica*, 50(6):1483–1504, 1982.

25 Subhash Khot, Richard J. Lipton, Evangelos Markakis, and Aranyak Mehta. Inapproximability results for combinatorial auctions with submodular utility functions. *Algorithmica*, 52(1):3–18, 2008.

26 Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior*, 55:270–296, 2006.

27 Takanori Maehara and Kazuo Murota. Valuated matroid-based algorithm for submodular welfare problem. *Annals of Operations Research*, 229:565–590, 2015.

28 Paul R. Milgrom. Putting auction theory to work: The simultaneous ascending auction. *Journal of Political Economy*, 108(2):245–272, 2000.

29 Paul R. Milgrom. The substitution metric and the performance of clock auctions, 2015. Talk at Simons Insitute, available at `https://simons.berkeley.edu/talks/paul-milgrom-10-13`.

30 Vahab S. Mirrokni, Michael Schapira, , and Jan Vondrák. Tight information-theoretic lower bounds for welfare maximization in combinatorial auctions. In *Proceedings of the 9th ACM Conference on Economics and Computation*, pages 70–77, 2008.

31 Kazuo Murota. Valuated matroid intersection I: Optimality criteria. *SIAM J. Discrete Math.*, 9(4):545–561, 1996.

32 Kazuo Murota. Valuated matroid intersection II: Algorithms. *SIAM J. Discrete Math.*, 9(4):562–576, 1996.

33 Noam Nisan and Ilya Segal. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory*, 129:192–224, 2006.

34 Michael Ostrovsky and Renato Paes Leme. Gross substitutes and endowed assignment valuations. *Theoretical Economics*, 2014.

35 J. G. Oxley. *Matroid Theory*. Oxford, 1992.

36 Renato Paes Leme. Gross substitutability: An algorithmic survey. Working paper, 2014.

37 Renato Paes Leme and Sam Chiu-wai Wong. Computing Walrasian equilibria: Fast algorithms and economic insights. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 632–651, 2017.

38 Christos H. Papadimitriou, Michael Schapira, and Yaron Singer. On the hardness of being truthful. In *Proceedings of the 49th Symposium on Foundations of Computer Science*, pages 250–259, 2008.

39 Christos H. Papadimitriou and Kenneth Steiglitz. *Combinatorial Optimization: Algorithms and Complexity*. Dover Publications, 2000.

40 Tim Roughgarden, Inbal Talgam-Cohen, and Jan Vondrák. When are welfare guarantees robust? In *2nd Algorithmic Game Theory and Data Science Workshop*, July 2016. Full version available at `https://arxiv.org/abs/1608.02402`.

41 A. Schrijver. *Combinatorial Optimziation: Polyhedra and Efficiency*. Springer, 2003.

42 Dravyansh Sharma, Amit Deshpande, and Ashish Kapoor. On greedy maximization of entropy. In *Proceedings of the 32nd International Conference on Machine Learning*, pages 1330–1338, 2015.

43 Yaron Singer and Jan Vondrák. Information-theoretic lower bounds for convex optimization with erroneous oracles. In *Proceedings of the 28th Neural Information Processing Systems Conference*, 2015.

44 Maxim Sviridenko, Jan Vondrák, and Justin Ward. Optimal approximation for submodular and supermodular optimization with bounded curvature. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1134–1148, 2015.

45 Rakesh V. Vohra. *Mechanism Design: A Linear Programming Approach*. Econometric Society Monographs, 2011.

46 Jan Vondrák. Optimal approximation for the submodular welfare problem in the value oracle model. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing*, pages 67–74, 2008.

A Standard Algorithms for Welfare-Maximization

Algorithms 1 and 2 describe the ascending price and the greedy approaches to welfare maximization.

ALGORITHM 1: The Kelso-Crawford ascending-price auction (formulated as an algorithm).

Input: Player valuations $v_1, \ldots, v_n$ represented by demand oracles; a parameter $\delta > 0$
Output: An allocation $\mathcal{S}$ and a price vector p
$p := 0$ and $\mathcal{S} := \emptyset$; % Initialization
while there exists a player i and a non-empty bundle D_i such that D_i is in demand given prices $p + \delta \vec{\mathbb{1}}_{j \notin S_i}$ and current allocation S_i, **do**
    $S_i := S_i \cup D_i$;
    $S_{i'} := S_{i'} \setminus D_i$ for every $i' \neq i$;
    $p(j) := p(j) + \delta$ for every $j \in D_i$;
end

ALGORITHM 2: Greedy maximization of value subject to a feasibility constraint.

Input: A valuation v represented by a value oracle; a feasibility constraint represented by a feasiblity oracle
Output: A bundle S
$S := \emptyset$; % Initialization
while there exists an item $j \notin S$ such that $S \cup \{j\}$ is feasible **do**
  Let j^* be an item that maximizes $v(j^* \mid S)$ among all items $j \notin S$ such that $S \cup \{j\}$ is feasible;
  $S := S \cup \{j^*\}$;
end

B Missing Proofs from Sections 3 and 4

Proof of Proposition 5. An additive valuation has the form $v(S) = \sum_{j\in S} a_j$. By assumption, $a_j \le v(\{j\}) \le (1+\epsilon)a_j$ for every singleton j. Hence we can determine the coefficients a_j within a factor of $1+\epsilon$ and run welfare maximization on the resulting additive functions. Each item simply goes to the highest bidder, and we lose a factor of at most $1+\epsilon$ due to the errors in determining a_j. ◀

Proof of Proposition 12. For every $S \subseteq T$ and $j \notin T$, let g_j be the decreasing function to which $v(j|\cdot)$ is ϵ-close. Then $(1+\epsilon)v(j|S) \ge (1+\epsilon)g_j(S) \ge (1+\epsilon)g_j(T) \ge v(j|T)$, showing that v is $(1+\epsilon)$-submodular. The converse follows from Observation 28. ◀

▶ **Observation 28.** *If a valuation v is α-submodular then it is marginal-ϵ-close for $\epsilon = \alpha - 1$ to the class of decreasing functions.*

Proof Sketch. Fix an item j. We want to show that the set function $v(j \mid \cdot)$ is ϵ-close to a decreasing function g_j. We define g_j as follows. For every bundle S, there is a range $[\frac{1}{1+\epsilon}v(j \mid S), v(j \mid S)]$ to which $g_j(S)$ must belong. Furthermore, for g_j to be decreasing, for every S and every subset $T \subseteq S$, the upper-bound on the range of $g_j(T)$ is also an upper-bound on the range of $g_j(S)$. We claim that: (1) If there is a non-zero range for $g_j(S)$ for every bundle S, then by picking the highest value in the range to be $g_j(S)$, we have defined a decreasing g_j such that $v(j \mid \cdot)$ is ϵ-close to it; (2) If there is a bundle S with negative range then we have found a contradiction to α-submodularity of v. Indeed, if this is the case then we have found a subset T such that $v(j \mid T) < v(j \mid S)/1+\epsilon = v(j \mid S)/\alpha$. This concludes the proof. ◀

Proof of Proposition 14. Define a linear valuation ℓ as follows: $\ell(S) = (1-c)(v(\emptyset) + \sum_{j=1}^{|S|} v(s_j \mid \emptyset))$ for every bundle S. Using that v has curvature c, $v(S) = v(\emptyset) + \sum_{j=1}^{|S|} v(s_j|S_j) \ge (1-c)(v(\emptyset) + \sum_{j=1}^{|S|} v(s_j|\emptyset)) = \ell(S)$. Using that v is α-submodular, $v(S) = v(\emptyset) + \sum_{j=1}^{|S|} v(s_j|S_j) \le v(\emptyset) + \alpha \sum_{j=1}^{|S|} v(s_j|\emptyset) \le \frac{\alpha}{1-c}\ell(S)$. Thus v is ϵ-close to ℓ for ϵ such that $1+\epsilon = \frac{\alpha}{1-c}$. ◀

Proof of Lemma 17.

$$\alpha \sum_{j=1}^{m_X} v(x_j \mid Y_j) \ \ge \ \sum_{j=1}^{m_X} v(x_j \mid Z \cup X_j) \tag{7}$$
$$= \ v(X \mid Z) \tag{8}$$
$$\le \ \ell(X \mid Z) - \epsilon\ell(Z), \tag{9}$$

where (7) follows since $Y_j \subseteq Z \cup X_j$ and since v is α-submodular, (8) follows by the disjointness of X and Z, and (9) follows since v is ϵ-close to ℓ and by Observation 9. ◀

C Kelso-Crawford and Close-to-Transversal Valuations

▶ **Lemma 29.** *In a market with α-submodular valuations $v_1, \ldots, v_n$ that are ϵ-close to unweighted transversal valuations $r_1, \ldots, r_n$, the Kelso-Crawford algorithm converges to a $\frac{1}{\alpha(1+3\epsilon)^2}$-biased Walrasian equilibrium.*

Proof. By monotonicity we can assume without loss of generality that all items are allocated at price 0 in the first step of the Kelso-Crawford algorithm, and once an item is allocated then it remains allocated throughout. Thus without loss of generality, Kelso-Crawford returns a full allocation as required by the definition of a μ-biased Walrasian equilibrium.

From now until the end of the proof, fix a player i. For simplicity we omit i from the notation. Let $\mathcal{P} = (P_1, \ldots, P_k)$ be the partition of the items corresponding to the unweighted transversal valuation r, to which the player's valuation v is ϵ-close. For an item j, let $P(j)$ denote the part to which this item belongs. We say that a part P is *represented* in a bundle X if X contains at least one item $j \in P$ with value $r(j) = 1$.

Towards proving the guarantee in Inequality (4), let S be the bundle allocated to the player by the Kelso-Crawford algorithm, and let T be an alternative bundle. Let $P(S)$ (resp., $P(T)$) be the parts represented in S (resp., T). Let $S' = \{j \in S \mid P(j) \in P(S) \cap P(T)\}$ be the set of items in S that belong to parts represented both in S and in T, and similarly $T' = \{j \in T \mid P(j) \in P(S) \cap P(T)\}$.

▶ **Claim 30.** $r(S') = r(T')$.

Proof of Claim 30. The value assigned to a set by valuation r is the number of parts represented in it. By definition, the same parts are represented in S' and in T'. ◀

We now relate the prices of S' and T' according to the price vector p with which the Kelso-Crawford algorithm terminated. In particular we show that the price of S' cannot be too high in comparison to the price of T'.

▶ **Claim 31.** $p(S') \leq p(T') + 3\alpha\epsilon r(S')$.

The proof of Claim 31 appears below. We use the above claims to complete the proof of Lemma 29. Let $S'' = S \setminus S'$ and $T'' = T \setminus T'$. We know that when the Kelso-Crawford algorithm terminates, the player cannot improve his utility by adding T'' to his allocation S, and so:

$$\begin{aligned} v(S) - p(S) &\geq v(S \cup T'') - p(S \cup T'') \\ &= v(S'') + v(S' \mid S'') + v(T'' \mid S) - p(S'') - p(S') - p(T''). \end{aligned} \tag{10}$$

By α-submodularity Lemma 25 applies and S'' is α-IR, therefore

$$v(S'') - p(S'') \geq (1 - \alpha)v(S''). \tag{11}$$

The marginal value assigned by r to a bundle X given a bundle Y is the number of parts represented in X but not in Y. Therefore $r(S' \mid S'') = r(S')$ and $r(T'' \mid S) = r(T'')$. By Observation 9 and since $r(S') = r(T')$ (Claim 30),

$$v(S' \mid S'') \geq r(S' \mid S'') - \epsilon r(S'') \geq r(S') - \epsilon r(S) = r(T') - \epsilon r(S); \tag{12}$$

$$v(T'' \mid S) \geq r(T'' \mid S) - \epsilon r(S) = r(T'') - \epsilon r(S). \tag{13}$$

Plugging (11)–(13) into (10), and using that $r(T')+r(T'') = r(T)$ and $p(S') \leq p(T')+3\alpha\epsilon r(S')$ (Claim 31),

$$\begin{aligned} v(S) - p(S) &\geq (1-\alpha)v(S'') + r(T) - 2\epsilon r(S) - p(T) - 3\alpha\epsilon r(S') \\ &\geq \frac{1}{1+\epsilon} v(T) - p(T) - (\alpha - 1 + 2\epsilon + 3\alpha\epsilon)v(S), \end{aligned}$$

where the last inequality uses v's ϵ-closeness to r. Thus $\mu \geq ((1+\epsilon)(\alpha+2\epsilon+3\alpha\epsilon))^{-1} \geq (\alpha(1+3\epsilon)^2)^{-1}$. This completes the proof of Lemma 29. ◀

Proof of Theorem 21. The proof follows directly from Proposition 23 combined with Lemma 29. ◀

C.1 Proof of Claim 31

Towards proving Claim 31, the following is a property of valuations that are α-submodular and ϵ-close to unweighted transversal valuations. It can be seen as a strengthening of Observation 9.

▶ **Claim 32.** *For every part P and bundles $X \subseteq P$ and Y, if P is represented in Y then $v(X \mid Y) \leq \alpha\epsilon$.*

Proof. Let y be an item representing P in Y. Then by α-submodularity, $v(X \mid Y) \leq \alpha v(X \mid y) \leq \alpha r(X \mid y) + \alpha\epsilon r(X \cup \{y\}) = \alpha\epsilon$, where the last inequality is by invoking Observation 9. ◀

Proof of Claim 31. Fix a part P that is represented in S' and T'. Let t be an item representing P in T'. Order the items representing P in S' according to the order in which the Kelso-Crawford algorithm added them to the player's allocation for the last time (i.e., at their termination prices according to p), breaking ties arbitrarily; denote the ordered set by $(s_1, s_2, \dots)$. Let B_j be the bundle of the player right before item s_j was added, and let D_j be the set of items (not including s_j) with which s_j was added to B_j.

We first consider item s_1. If it is the case that P is not represented in B_1 and no other item represents P in D_1, then we argue that $v(s_1 \mid B_1 \cup D_1) - p(s_1) \geq v(t \mid B_1 \cup D_1) - p(t)$. The reason for this is that otherwise, the player's utility could have been improved by replacing s_1 by t (using that t's termination price $p(t)$ is weakly higher than its price when s_1 was added). By monotonicity, we can write $v(s_1 \mid B_1 \cup D_1) \leq v(s_1, t \mid B_1 \cup D_1) \leq v(t \mid B_1 \cup D_1) + \alpha\epsilon$, where the last inequality is by Claim 32. Therefore $p(s_1) \leq p(t) + \alpha\epsilon$. In the remaining case, P is represented in either B_1 or D_1. We know that $v(s_1 \mid B_1 \cup D_1) \geq p(s_1)$, otherwise the utility could have been improved by dropping s_1. By Claim 32, the left-hand side is at most $\alpha\epsilon$, and so $p(s_1) \leq \alpha\epsilon$.

Now consider the rest of the items $s_2, s_3, \dots$. When item s_j is added, $\{s_1, \dots, s_{j-1}\} \subseteq B_j \cup D_j$. We also know that, as above, $v(s_j \mid B_j \cup D_j) \geq p(s_j)$. By α-submodularity this means that $\alpha \sum_{j\geq 2} v(s_j \mid S_j) \geq \sum_{j \geq 2} p(s_j)$. The left-hand side is equal to $\alpha v(\{s_2, s_3, \dots\} \mid s_1)$, and is therefore $\leq \alpha\epsilon$ by Observation 9. We have thus shown that $\sum_{j\geq 2} p(s_j) \leq \alpha\epsilon$, and the same argument shows that the total payment for items in $S' \cap P$ that do not represent P is at most $\alpha\epsilon$.

We conclude that the total payment for items in $S' \cap P$ is at most $p(t) + 3\alpha\epsilon$. Summing up over all parts represented in S' and T' completes the proof of the claim. ◀

D Negative Results for Specific Algorithms

How well do standard algorithms for welfare maximization perform for valuations that are close to, but are not quite, gross substitutes? After discussing the LP-based approach in Section 3.2, in this section we discuss the two other main approaches to welfare maximization for gross substitutes – the ascending auction algorithm of Kelso and Crawford [24], and the cycle canceling algorithm of Murota [31, 32]. Some of the details are deferred to the full version [40] due to space limitations.

At first glance, the Kelso-Crawford algorithm seems like a promising approach due to its known welfare guarantee of a 1/2-approximation for the class of submodular valuations [17]. However Proposition 20 and its proof in the full version [40, Proposition 8] show that the Kelso-Crawford algorithm cannot in general guarantee much better than that, even for simple submodular valuations that are arbitrarily close to unit-demand. More precisely, for every $\epsilon \leq 1$ there exists a market with $O(1/\epsilon)$ players whose submodular valuations are ϵ-close to unit-demand, and for which the Kelso-Crawford algorithm with adversarial ordering of the players finds an allocation with $\approx 2/3$ of the optimal welfare. An interesting open question is whether the Kelso-Crawford algorithm can be modified to eliminate such bad examples, e.g., by ordering the players in an optimal way.

The cycle canceling algorithm of Murota is a different approach which relies on properties of GS (or equivalently $M^\natural$-concave) valuations under local improvements. In Section D.1 we show that such properties hold for submodular valuations in a certain approximate sense, but unfortunately the cycle canceling algorithm cannot find a local optimum that would allow us to exploit these properties. In fact we show that a local optimum in Murota's sense does not exist for submodular valuations.

D.1 Murota's Cycle Canceling Approach

The cycle canceling approach is based on the following useful property of gross substitutes valuations, called the *single improvement (SI)* property [18]: Given a price vector p, we say that a bundle S is in *local demand* for a valuation v if its utility cannot be improved by adding an item, removing an item or swapping one item for another. The SI property holds if every bundle S in local demand is also in (global) demand (i.e., $v(S) - p(S) \geq v(T) - p(T)$ for every bundle T). Murota's cycle canceling algorithm finds an allocation and prices such that each player gets a bundle in local demand, thus arriving at an optimal allocation for gross substitutes valuations.

We observe that a variant of the SI property characterizes submodular valuations, and thus the SI property interpolates between submodularity and gross substitutes.

▶ **Definition 33.** Let $\beta \in [0, 1]$. A valuation v is β-SI if for every S in local demand, S is strongly individually rational,[14] and $v(S) - \beta p(S) \geq v(T) - p(T)$ for every T.

▶ **Observation 34.** *A monotone valuation is submodular if and only if it is* 0*-SI. A monotone valuation is gross substitutes if and only if it is* 1*-SI.*

Proof. Suppose v is submodular; we want to prove that v is 0-SI. Let S be in local demand, $v(S + i) - v(S) \leq p_i$, and $v(S - j) - v(S) \leq -p_j$. (We don't even need the swap property.) By submodularity and the second property, we have $v(S \cup T) \leq v(S) + \sum_{i \in T \setminus S} p_i$. Therefore,

[14] Being strongly IR implicitly holds for SI as well, i.e., if v is SI and S is in local demand then S is strongly IR.

$v(S) \geq v(S \cup T) - p(T \setminus S) \geq v(T) - p(T)$ by monotonicity. Also, for every $S' \subset S$, by submodularity we have $v(S \setminus S') \leq v(S) - \sum_{j \in S'} p_j$. Therefore $v(S') \geq v(S) - v(S \setminus S') \geq p(S')$.

Conversely, suppose that v is not submodular, i.e. $v(S+a+b) > v(S) + v(a \mid S) + v(b \mid S)$ for some S and $a, b \notin S$. We set $p_i = 0$ for $i \in S$ and $p_j = v(j \mid S)$ for $j \notin S$. Clearly S is in local demand, because swapping at most 1 element does not increase utility. However, $v(S \cup \{a, b\}) - p(S \cup \{a, b\}) > v(S)$, so v is not 0-SI. ◀

Moreover, as the next observation shows, finding an allocation with prices such that each player gets a bundle in their local demand would give a smooth transition would provide a smooth transition between a 1/2-approximation for submodular valuations and an optimal solution for gross substitutes.

▶ **Observation 35.** *For β-SI valuations, any allocation (of all items) with prices such that each player gets a bundle in their local demand has value at least $\frac{1}{2-\beta}\,\mathrm{OPT}$.*

Proof. Suppose that each v_i is β-SI and we have an allocation $(A_1, \ldots, A_n)$ with prices such that each A_i is in the local demand of player i. We assume that all items are allocated, and the allocation is individually rational; therefore, $\sum_{i=1}^n v_i(A_i) \geq \sum_{i=1}^n p(A_i) = \sum p_j$.

Suppose that the optimal allocation is $(O_1, \ldots, O_n)$. By the β-SI property, we have $v_i(A_i) - \beta p(A_i) \geq v_i(O_i) - p(O_i)$. Adding up over all players, $\sum_{i=1}^n v_i(A_i) - \beta \sum p_j \geq \mathrm{OPT} - \sum p_j$. From here, $\sum_{i=1}^n v_i(A_i) \geq \mathrm{OPT} - (1-\beta) \sum p_j \geq \mathrm{OPT} - (1-\beta) \sum_{i=1}^n v_i(A_i)$. Thus $(2-\beta) \sum_{i=1}^n v_i(A_i) \geq \mathrm{OPT}$. ◀

Unfortunately, this approach (whether by cycle canceling or by any other method) is doomed to fail: Example 1 in the full version [40] shows a market with submodular valuations for which it is impossible to find an allocation and prices such that every player's bundle is in their local demand. In other words, for this market it is impossible to get rid of all negative cycles in Murota's algorithm, ruling out this approach.