

Algebraic Markets: Robust Equilibrium Existence

ARTHUR DOLGOPOLOV

Department of Economics and Finance, University of Rome Tor Vergata

I demonstrate an algebraic framework for markets with indivisible goods where competitive equilibrium prices are guaranteed to exist for every admissible profile of preferences if and only if the corresponding welfare-maximization problem can be solved efficiently. In other words, competitive markets solve precisely the class of all tractable optimization problems. This result holds for any valuation class that satisfies the algebraic assumption and applies to markets involving both private and public goods. The findings build on a novel connection between market equilibria and valued constraint satisfaction problems from complexity theory. I illustrate how this approach can be employed to establish the existence or nonexistence of competitive prices in several new settings.

KEYWORDS. Asymptotic complexity, Competitive equilibrium, Valued constraint satisfaction, Supermodularity.

JEL CLASSIFICATION. First JEL, C62, D47, C61.

1. INTRODUCTION

When does a market equilibrium provably exist? A well-known sufficient condition is the gross substitutes property (Kelso Jr and Crawford (1982), Gul and

Arthur Dolgoplov: Dolgoplov@Economia.uniroma2.it

I thank Elizabeth Baldwin, Juha Tolvanen, Lasse Mononen, Marcus Pivato, and Marek Pycia, anonymous referees at EC'2026, and audiences at Bielefeld University and the Workshop on Collective Decisions for their insightful comments on the early versions of this paper.

Stacchetti (1999)), but it is often too restrictive for practical market design applications (Brandenburg et al. (2023)). This paper proposes a new, complexity-theoretic approach to identifying classes of valuations that guarantee equilibrium existence.

An instructive example comes from Bikhchandani and Ostroy (2002), who showed that in a particular combinatorial auction environment of a package assignment problem, a market equilibrium exists if and only if a specific integer program can be solved via the corresponding (non-integer) linear programming approximation. Remarkably, constraint matrices of bipartite assignment problems have the total unimodularity property¹, which renders the solution of the linear program integral (Hoffman (1956)). Hence, for this class of problems, equilibria always exist. Since linear programming is computationally easy, so is the problem of finding efficient allocations for assignment problems, which was already noted by Von Neumann (1950). This raises the question of whether the tight link between polynomial-time solvability and equilibrium existence extends beyond this special case.

To this end, the present paper makes two contributions. On the theoretical side, it shows that, within a particular algebraic framework of markets over discrete goods, competitive equilibria exist exactly for those classes of allocation problems that are tractable in the algorithmic sense, that is, solvable in polynomial time. The implication from equilibrium existence to tractability is intuitive, since computing an equilibrium yields an efficient allocation. The converse — that any tractable problem corresponds to a market that admits a competitive equilibrium — is less immediate and is the main conceptual result. When equilibria exist, the supporting prices can also be found in polynomial time from the dual program.

On the practical side, the paper develops a method for identifying classes of valuations and allocation problems for which equilibrium existence is guaranteed. The party organizing a combinatorial auction or an exchange (the ‘designer’) must fix an elicitation and pricing format ex ante, before observing bids:

¹A matrix is totally unimodular if every square submatrix has determinant 0, 1 or -1 .

for example, permitting only item prices, or allowing limited forms of package bids. The designer has to know for which classes of valuations this format always admits prices that support an efficient allocation in equilibrium, to guarantee convergence to an efficient allocation. Complexity of finding the efficient allocation can be used as an indicator of whether competition will reach this outcome, and can be mined for new positive or negative results: classes of valuations that guarantee existence or imply non-existence. Consider a market with pairwise complementarities: a number of broadcasters are competing for consumers living on an undirected graph. The broadcasters can sometimes economize by providing the service to a pair of neighboring subscribers. It is known from theoretical computer science that finding an efficient assignment of subscribers to broadcasters can be done efficiently for 2 broadcasters, and cannot be done for 3 or more ([Gridchyn and Kolmogorov \(2013\)](#)). This paper's approach shows that competitive equilibria are also guaranteed if and only if there are 2 broadcasters.

To further illustrate the approach, I additionally assume that buyers are allowed to bid, at a minimum, on each individual good (and possibly also on bundles). In auction terms, this is equivalent to requiring that singleton bids are admissible. Under this assumption, the classes of valuations that guarantee equilibrium existence are easily characterized: equilibrium exists for all valuations in the class if and only if they are supermodular. This has a stark negative implication for market design: if additively separable utility functions cannot reasonably be ruled out, then robust equilibrium existence can be guaranteed only in the narrow case of supermodular valuations. For any other class of valuations, there exists a profile of buyer preferences for which no competitive equilibrium exists.

Finally, from the connection to linear programming, I derive a generic, auction-like iterative procedure that converges to a competitive equilibrium whenever one exists. As in ascending-price mechanisms ([Parkes \(2005\)](#), [Mishra and Parkes \(2007\)](#), [Blumrosen and Nisan \(2010\)](#)), the procedure is driven by 'demand queries': at each step, buyers and sellers report their payoff-maximizing choices at current prices, and prices are then updated. Under truthful reporting, this eliminates the need for the designer to know agents' complete preferences in order to compute the equilibrium outcome.

The proof of the main equivalence combines two existing ideas: the finite-valued VCSP dichotomy of [Thapper and Živný \(2016\)](#), which characterizes when the basic linear programming relaxation solves a valued constraint language, and the linear-programming duality argument behind equilibrium existence in package-assignment environments, as in [Bikhchandani et al. \(2001\)](#). The contribution is the identification of a market equilibrium object for which these two results become exactly equivalent. This turns complexity classifications into equilibrium-existence and non-existence results for whole classes of markets, including non-private-good environments. Further results extend the connection to allow for more flexibility in buyer and seller utilities than direct transfer of the theorems would imply.

Throughout, agents have quasilinear preferences in a numeraire and face no binding budget constraints, as in [Kelso Jr and Crawford \(1982\)](#) and treatments of Walrasian equilibrium in market design. In particular, every agent is assumed to have enough money to afford any bundle that arises in equilibrium, so that prices serve purely as transfers rather than feasibility constraints. The setup includes standard anonymous item prices as a special case in private-goods assignment environments, but it also includes Lindahl-style personalized prices, which are natural in public-good and externality environments. The model is also conceptually similar to a version of the Arrow–Debreu model of general equilibrium, known as the generalized Nash equilibrium problem (GNEP) ([Facchinei and Kanzow \(2010\)](#)), but with key differences: the discrete nature of choices and the inclusion of a numeraire good.

Existing work already uses linear programming to test equilibrium existence at the level of a specific auction instance. In the package assignment model of [Bikhchandani and Ostroy \(2002\)](#), the designer can check ex post whether there exist competitive prices that support the submitted bids. The present paper is concerned instead with imposing ex ante restrictions on the allowable bids or valuations that guarantee robust equilibrium existence at the level of a whole class of markets. The difference is illustrated by [Baldwin and Klemperer \(2019\)](#), who identify unimodularity as a structural condition that guarantees class-level equilibrium existence, and relate it to a different instance-level condition, the

Intersection Count Theorem. A different route is taken here, which frames the search for efficient allocations in terms of computational complexity and proposes a complexity-theoretic criterion for when equilibria are guaranteed to exist. It is also worth noting that neither framework nests the other. [Baldwin and Klemperer \(2019\)](#) develop a theory for private-goods environments with bundle demand and structural restrictions on demand types. By contrast, the present paper generalizes to allocation problems that need not be private-goods markets, but does not allow for structural restrictions. This permits a Lindahl-style interpretation of equilibrium prices in collective-choice environments with quasilinear utility, such as [Delacrétaz et al. \(2025\)](#). It also accommodates algorithmically nontrivial structured non-private-good environments, such as the technological dependencies setting in Example 3, which, to the best of our knowledge, has not previously been studied from the viewpoint of equilibrium existence or computation. At the same time, it excludes known classes with robust equilibrium existence such as OR-of-XORs-of-singletons ([Lehmann et al. \(2006\)](#)) precisely because structural restrictions would be necessary. Proposition 2 confirms this as a general negative result: without structurally restricting utility functions, the possibilities for robust equilibrium existence are extremely narrow.

A closely related work by [Roughgarden and Talgam-Cohen \(2015\)](#) identifies a necessary complexity-theoretic condition for robust equilibrium existence, in terms of a separation between the difficulty of consumers' optimization problems and that of market-wide optimization. The results here can be viewed as a converse: by imposing more structure on the market classes, I show that some market in any intractable class must fail to admit a competitive equilibrium.

There are also several tractable cases in which equilibrium existence is known, including nonlinear and non-anonymous pricing ([Bikhchandani and Ostroy \(2002\)](#)), graphical pricing ([Brandenburg et al. \(2023\)](#)), gross substitutes ([Gul and Stacchetti \(1999\)](#)), unit bids ([Lehmann et al. \(2006\)](#)), geometric structures ([Rothkopf et al. \(1998\)](#), [Candogan et al. \(2015\)](#)), and certain tractable bidding languages ([Nisan \(2000\)](#), [Nisan and Segal \(2006\)](#)). These conditions are typically sufficient but not necessary. The present characterization results suggest that such examples are a part of a broader complexity-theoretic picture.

The remainder of the paper is organized as follows. Section 2 introduces the abstract market framework for allocating discrete goods and defines classes of markets. Section 3 states the main theorem, sketches the idea of the proof, and illustrates its implications through examples. Section 4 refines the results for the particular case of conservative markets. Sections 5 and 6 develop applications: the former shows how the main algebraic assumption can be relaxed when buyers agree on the structure of the market, and the latter offers an iterative price-adjustment procedure that arrives at a market equilibrium. Section 7 discusses the possibility of allowing for infeasible allocations and limitations of the framework. The concluding section discusses possible further relaxations of the algebraic assumption. All proofs are collected in the [Appendix](#).

2. DEFINITIONS

2.1 Markets

Let Y be a finite set of tradable goods and let B be a finite set of buyers. Let D be a finite set, called the domain. An allocation is a function $\omega : Y \rightarrow D$, and the set of all allocations is

$$\Omega = \{\omega : Y \rightarrow D\}.$$

The abstraction of the domain D allows a unified treatment of several environments. If $D = B$, then an allocation assigns each good to a buyer. If $D = \{0, 1\}$, then an allocation selects which goods are provided as public goods (those mapped to 1).

Each buyer $b \in B$ has a rational-valued utility function over allocations, $U_b : \Omega \rightarrow \mathbb{Q}$. Each good $y \in Y$ has an associated seller whose utility depends only on the chosen domain value for that good, $V_y : D \rightarrow \mathbb{Q}$, which captures (in particular) production costs.

Define utilitarian welfare in the usual way as

$$W(\omega) = \sum_{b \in B} U_b(\omega) + \sum_{y \in Y} V_y(\omega(y)).$$

These primitives already define the welfare-maximization problem $\max_{\omega \in \Omega} W(\omega)$ without the need for prices or equilibrium concepts, which will be introduced later.

DEFINITION 1 (Market). A *market* is a tuple

$$M = \langle B, Y, D, (U_b)_{b \in B}, (V_y)_{y \in Y} \rangle,$$

where B is a finite set of buyers, Y is a finite set of goods, D is a finite domain, $\Omega = \{\omega : Y \rightarrow D\}$, each $U_b : \Omega \rightarrow \mathbb{Q}$, and each $V_y : D \rightarrow \mathbb{Q}$.

Together with the objective $\max_{\omega \in \Omega} W(\omega)$, a market induces an unconstrained optimization problem. To discuss its complexity, one will need to define how the functions U_b and V_y are represented, which is postponed until the next section.

Definition 1 extends the standard Walrasian equilibrium by allowing buyers to pay to assign individual goods to elements in the domain D . Prices are allowed to depend on the buyer as well as on the chosen domain value for each good. The revenue goes to the sellers of the respective goods.

A *price system* is a function $p : Y \times D \times B \rightarrow \mathbb{Q}$. The value $p(y, d, b)$ is interpreted as the transfer from buyer b to the seller of good y if good y is assigned value $d \in D$. This could be the payment for selling a private good to buyer b (if $D = B$) or for providing a public good (if $D = \{0, 1\}$).

Buyers have quasilinear payoffs and face no binding budget constraints. Given prices p , buyer b 's demand correspondence is

$$\mathcal{D}_b(p) = \arg \max_{\omega \in \Omega} \left(U_b(\omega) - \sum_{y \in Y} p(y, \omega(y), b) \right).$$

Similarly, the seller of good y chooses a domain value $d \in D$ to maximize quasilinear payoff, yielding the supply correspondence

$$\mathcal{S}_y(p) = \arg \max_{d \in D} \left(V_y(d) + \sum_{b \in B} p(y, d, b) \right).$$

The maximized value for the seller is the sum of all transfers from buyers to the seller and the utility function of the seller, which captures the costs of production.

If $D = B$, then the seller chooses the buyer that would pay the highest price for y . If $D = \{0, 1\}$, then the problem instead becomes the decision to provide the public good or not.

DEFINITION 2. A *market equilibrium* for a market $M = \langle B, Y, D, (U_b)_{b \in B}, (V_y)_{y \in Y} \rangle$ is a pair (ω, p) with $\omega \in \Omega$ and $p : Y \times D \times B \rightarrow \mathbb{Q}$ such that:

- (i) **Buyer optimality:** For every $b \in B$, $\omega \in \mathcal{D}_b(p)$.
- (ii) **Seller optimality:** For every $y \in Y$, $\omega(y) \in \mathcal{S}_y(p)$.

The two equilibrium conditions say that at prices p , no buyer wants to switch to a different allocation and no seller wants to change the chosen allocation of their good.

The definition warrants some discussion. First, the price system is less permissive than it may seem from the triplet of inputs. Although prices are allowed to depend on both buyers and goods, the price of a good does not depend on a bundle chosen by the buyer, or on the assignment of other goods. In this sense the price system is linear in the allocation coordinates. The price system is, however, generally non-anonymous: different buyers may face different transfers for the same pair y, d . For private-goods environments, where $D = B$, anonymous Walrasian item prices are obtained as the special case. Since the sellers are maximizing their payoff by choosing an element $d \in D$, only the highest price among all such d will constitute a transfer. The lower ‘prices’ for other buyers are then the counterfactual transfers, i.e. unsuccessful bids or their willingness to pay, which will not be paid in equilibrium. The equilibrium conditions can be equivalently interpreted as an auction with a single seller or an exchange with many independent sellers, because every seller’s payoff is independent of the choices of the other sellers. In addition to Walrasian anonymous prices, the setup also includes non-anonymous pricing of public goods.

Prices are only defined up to additive constants that do not affect optimal choices. In particular, fixing any $b \in B$ and $y \in Y$, adding the same constant to all values $\{p(y, d, b)\}_{d \in D}$ shifts payoffs but does not change the solutions to the buyers’ and sellers’ $\arg \max$ problems.

REMARK 1 (Additive normalization). If (ω, p) is a market equilibrium and $b' \in B$, $y' \in Y$, $c \in \mathbb{Q}$, define p' by

$$p'(y', d, b') = p(y', d, b') + c \quad \text{for all } d \in D, \quad p'(y, d, b) = p(y, d, b) \text{ otherwise.}$$

Then (ω, p') is also a market equilibrium.

Therefore, prices remain agnostic about the endowments of players. It is always possible to pick for every $b \in B$ and allocation ω^e prices p' , such that $p'(y, d, b) = p(y, d, b) - p(y, \omega^e(y), b)$. Facing prices p' , buyer b always has an option to demand ω^e for zero price. In classical terms, this allocation is the endowment, and this property is the standard requirement of individual rationality: a choice where no money changes hands is always available to every buyer and every seller.

2.2 Classes of markets

In practice, the exact utility functions of the buyers are likely to be unknown. Instead, it may be accepted that they belong to a certain class, and the goal is to ensure that the market equilibrium robustly exists for all possible markets in that class. For example, an organizer of an auction may know that buyers view certain goods as substitutes, but does not know ex ante the identity of substitute goods nor the exact utility functions of the consumers, including the magnitude of the substitution effect. To model such classes of preferences, I use market classes, which are defined below. First, I define scopes and subutility functions to separate the structural restrictions on preferences from the specific goods involved.

A scope x is an ordered tuple of goods $x = (x(1), x(2), \dots, x(k))$, where each $x_i \in Y$. To save on notation, I will sometimes write x equivalently as an injective mapping $x : \{1, 2, \dots, k\} \rightarrow Y$, and thus $x(r)$ for the r -th element of x . In this case $y \in x$ always means $y \in \text{Im}(x)$. In particular, with some further abuse of notation, let $\omega(x)$ denote componentwise composition of ω and x , i.e. the vector $(\omega(x(i)))_{1 \leq i \leq k}$. Practically, this is a restriction of allocation ω to the goods in x . Let D^x be the set of all such mappings for a fixed scope x , $D^x = \{d : x \rightarrow D\}$. A scope is a generalization of a package in auction theory: whereas a package is typically




FIGURE 1. Scopes and subutility functions.

an unordered set of goods, a scope is ordered to map precisely to the arguments of the subutility function.

A subutility function specifies a utility level for different partial allocations of goods in a given scope. Let $\mathcal{F}_k$ denote the set of all functions $f : D^k \rightarrow \mathbb{Q}$ for some finite $k \geq 1$, and let $\mathcal{F} = \bigcup_k \mathcal{F}_k$. For any function $f \in \mathcal{F}_k$, k is the arity of f , which is written $ar(f) = k$.

A component is a pair (x, f) , where x is a scope and $f \in \mathcal{F}_k$ is a subutility function with arity k equal to the length of x . Components encode nonlinear relationships in the buyers' utility functions while maintaining additive separability across components. The scope and the subutility in each component separate the structure from the functional form. The subutility function f defines the functional form of the relationship (e.g., complementarity), while the scope x specifies which particular goods are subject to this relationship.

To illustrate the concepts of scopes and subutility functions, let $D = B$ and consider a market with five goods $Y = \{y_1, \dots, y_5\}$ shown in Figure 1a. Three pairs of goods are complements: if both goods in such a pair are assigned to the same buyer, that buyer's utility increases by 1. This relationship is illustrated by the weighted edges and is captured by the subutility function $\sigma^1 \in \mathcal{F}_2$, defined by

$$\sigma^1(a_1, a_2) = \begin{cases} 1 & \text{if } a_1 = a_2, \\ 0 & \text{otherwise.} \end{cases}$$

The value is positive only when the two goods are assigned the same domain value (interpreted as both goods being allocated to the same buyer). However,

the arguments a_1 and a_2 may correspond to any pair of goods in Y . Thus, a market in $\text{Market}(\{\sigma^1\})$ might include a component $((y_1, y_4), \sigma^1)$, or $((y_2, y_3), \sigma^1)$, or both as in Figure 1a. Figure 1b shows a different market: the subutility function is the same, but it is applied to different scopes. Finally, Figure 1c shows the converse: the scopes are the same as in 1a, but the subutility differs (substitutability rather than complementarity), so σ^1 no longer represents the interaction.

Each finite subset $\mathcal{U} \subseteq \mathcal{F}$ defines a market class, denoted $\text{Market}(\mathcal{U})$.

DEFINITION 3 (Market class). Market $\langle B, Y, D, (U_b)_{b \in B}, (V_y)_{y \in Y} \rangle$ belongs to the market class $\text{Market}(\mathcal{U})$ if

- (i) $V_y \in \mathcal{U}$ for any $y \in Y$ and
- (ii) for every $b \in B$, there exists a finite set of components $\{(x_{b,t}, f_{b,t})\}_{t \in T_b}$ such that $f_{b,t} \in \mathcal{U}$ for all t , and

$$U_b(\omega) = \sum_{t \in T_b} f_{b,t}(\omega(x_{b,t})) \quad \text{for all } \omega \in \Omega.$$

Since the class restricts only the permissible subutility functions and not the scopes, it describes possible types of relationships between goods without mandating which specific goods enter these relationships or restricting the total number of goods.

Any market in a class $\text{Market}(\mathcal{U})$ induces a corresponding unconstrained optimization problem of maximizing utilitarian welfare, where individual utility functions of buyers are expressed through function–package tuples:

DEFINITION 4 (Optimization problem). The optimization problem for a market in $\text{Market}(\mathcal{U})$ is to find an allocation $\omega \in \Omega$ that maximizes the objective function $W : \Omega \rightarrow \mathbb{Q}$ defined by:

$$W(\omega) = \sum_{b \in B} U_b(\omega) + \sum_{y \in Y} V_y(\omega(y)) = \sum_{b \in B} \sum_{t \in T_b} f_{b,t}(\omega(x_{b,t})) + \sum_{y \in Y} V_y(\omega(y)).$$

The solution of this problem is an optimal allocation. Polynomial-time computability can now be defined for the whole market class: $\text{Market}(\mathcal{U})$ is said to

be tractable if for any finite $\mathcal{U}' \subseteq \mathcal{U}$ there is an algorithm² that computes an optimal allocation for any market in $\text{Market}(\mathcal{U}')$ in time polynomial in the number of goods $|Y|$.

A consequence of this algebraic setup is that the market class stores no information on scopes. The set $\mathcal{U}$ defines only the shapes of allowable subutility functions, but does not declare which goods they will be applied to. For example, the class $\text{Market}(\{\sigma^1\})$ guarantees only that buyers *may* have complementary preferences over some pairs of goods. This property is formalized by the following remark:

REMARK 2 (Algebraicity). Let $M = \langle B, Y, D, (U_b)_{b \in B}, (V_y)_{y \in Y} \rangle$ be a market in the class $\text{Market}(\mathcal{U})$. For any $\omega \in \Omega$ and any $b \in B$, let U'_b be such that $U'_b(\omega) = U_b(\omega_b)$, where $\omega_b(y) = \omega(\rho_b(y))$ for all $y \in Y$ and some permutation $\rho_b : Y \rightarrow Y$. Then M is also in $\text{Market}(\mathcal{U})$.

Algebraicity says that if a valuation structure is permissible for a certain set of goods, it is permissible for any permutation of those goods. It is a restriction on market classes, not on individual market instances. In a given instance, the scopes may encode any pattern of relationships among named goods. For example, one pair of goods may enter a complementary component, another pair may enter a substitutable component, and a third pair may not enter any component at all. What algebraicity rules out is not such instance-level heterogeneity, but rather class-level restrictions on which scopes are admissible. Thus $\text{Market}(\mathcal{U})$ may contain markets with arbitrary complementarity graphs, but it cannot define a class by requiring that the complementarity graph be bipartite, planar, etc.

²The input of such an algorithm would have to include a representation of all functions in $\mathcal{U}'$. From the perspective of asymptotic complexity, the specific representation does not matter. For example, one could provide a complete table of values for every input to every function.

3. MAIN RESULT

The main result of this paper is an if-and-only-if connection between tractability of welfare maximization and the existence of market equilibria.³

PROPOSITION 1. *Every market in the class $\text{Market}(\mathcal{U})$ has a market equilibrium if and only if every corresponding optimization problem can be solved in time polynomial in the number of goods. In this case, every efficient allocation can be supported by a market equilibrium.*

Proposition 1 suggests a practical test for equilibrium existence in a given market. If the market can be formulated using only functions from a tractable class, then a market equilibrium exists. This can often be checked using existing dichotomy results and Section 4 illustrates this approach.

The proof of Proposition 1 is constructive and yields slightly more than equilibrium existence. Whenever the welfare-maximization problem for a market in $\text{Market}(\mathcal{U})$ is solved by a linear program, the dual program directly provides supporting equilibrium prices. Hence, in every tractable class covered by Proposition 1, one can compute not only an efficient allocation but also a competitive equilibrium (ω, p) in polynomial time.

COROLLARY 1. *If a market M belongs to a tractable class $\text{Market}(\mathcal{U})$, then M admits a market equilibrium.*

The converse is not true. A tractable market instance may belong to both tractable and intractable classes. For example, markets with completely separable utility functions clearly admit market equilibria. However, the class of such markets is a subset of larger, intractable classes—such as the conservative market classes discussed in Section 4—which also contain markets that lack equilibria. Thus, while membership in a tractable class guarantees existence, membership in an intractable class does not preclude it.

³We make the standard assumption throughout that $P \neq NP$, otherwise the results are vacuous

3.1 Proof idea and constraint satisfaction

The proof of Proposition 1 relies on the concept of a valued constraint satisfaction problem (VCSP) and its properties.

DEFINITION 5. An instance of a finite-valued constraint satisfaction problem (VCSP) is a triple $\langle Y, D, f \rangle$ where:

- Y is a finite set of variables,
- D is a finite set, called the domain,
- f is the objective function, mapping assignments of variables $\Omega = \{\omega : Y \rightarrow D\}$ to $\mathbb{Q}$. It is given by

$$f(\omega) = \sum_{t \in T} f_t(\omega(x_t)),$$

with T standing for a finite set. Function f is given by a set of components (x_t, f_t) , consisting of a function $f_t \in \mathcal{F}_{k_t}$ and a scope $x_t = (x_{t,1}, \dots, x_{t,k_t})$, where every $x_{t,i} \in Y$.

Since there is a clear correspondence between any market $\langle B, Y, D, (U_b)_{b \in B}, (V_y)_{y \in Y} \rangle$ and a VCSP instance of maximizing the total welfare of this market $\langle Y, D, W(\omega) = \sum_{b \in B} U_b(\omega) + \sum_{y \in Y} V_y(\omega) \rangle$, same notation is used for both. As before, $\omega(x_t)$ denotes the restriction of assignment ω to the variables in x_t .

Such VCSPs are called finite-valued because the range of all constraint functions f_t is $\mathbb{Q}$ (finite values), rather than $\mathbb{Q} \cup \{-\infty\}$. In market context, this corresponds to the assumption that all allocations are feasible. A solution to an instance of VCSP is an assignment $\omega : Y \rightarrow D$ that maximizes the objective function f .⁴

VCSP classes are defined in terms of allowed constraint functions. A valued constraint language $\Gamma \subseteq \mathcal{F}$ is a set of functions, and $\text{VCSP}(\Gamma)$ denotes the class

⁴In order to capture welfare maximization, the definition uses a more suggestive formulation as a maximization problem rather than minimization that is standard in the constraint satisfaction literature.

of all VCSP instances in which every function f_t belongs to Γ . Analogous to market classes, the functions in Γ do not encode information about which variables they apply to, only instances of VCSP specify the scopes. For example, if $\Gamma = \{f_1\}$ where $f_1 \in \mathcal{F}_2$, an instance with variables $Y = \{y_1, y_2, y_3, y_4\}$ might contain constraints like $f_1(\omega(y_1, y_2))$, $f_1(\omega(y_3, y_4))$, or $f_1(\omega(y_2, y_4))$.⁵

A valued constraint language Γ is said to be tractable if for every finite subset $\Gamma' \subseteq \Gamma$, any instance in $\text{VCSP}(\Gamma')$ can be solved in polynomial time. Otherwise, it is *intractable*. A key tool for analyzing VCSP complexity is a specific algorithm called the basic linear programming relaxation (BLP) of a VCSP instance. It is defined as the following linear program:

$$\begin{aligned}
 & \max \sum_{t \in T} \sum_{\mathbf{d} \in D^{x_t}} f_t(\mathbf{d}) \lambda_{t,\mathbf{d}}, \\
 & \text{s.t.} \quad \sum_{\substack{\mathbf{d} \in D^{x_t}: \\ \mathbf{d}(r)=d}} \lambda_{t,\mathbf{d}} = \mu_{x_t(r),d} \quad \forall t \in T, d \in D, r \in \{1, \dots, k_t\}, \\
 & \quad \sum_{d \in D} \mu_{y,d} = 1 \quad \forall y \in Y, \\
 & \quad \sum_{\mathbf{d} \in D^{x_t}} \lambda_{t,\mathbf{d}} = 1 \quad \forall t \in T, \\
 & \quad 0 \leq \lambda, \mu \leq 1.
 \end{aligned} \tag{1}$$

The variables in this linear program can be interpreted as probability distributions. Namely, $\mu_{y,d}$ is the probability that variable $y \in Y$ is assigned value $d \in D$, and $\lambda_{t,\mathbf{d}}$ is the probability that the variables in the scope x_t of constraint t are assigned the tuple of values $\mathbf{d} \in D^{x_t}$. The first set of constraints ensures that the marginal probability of a variable y taking value d must match the probability implied by $\lambda_{t,\mathbf{d}}$ for any constraint t that involves y . Specifically, if variable y appears at the r -th position in scope x_t (i.e., $x_t(r) = y$), the sum of $\lambda_{t,\mathbf{d}}$ over all tuples $\mathbf{d}$

⁵Such constraints are called relational in VCSP literature. The counterpart, structural constraints that restrict the scopes and compositions of functions, are an active field of study. Combinations of the two types of constraints are called hybrid constraints [Cooper and Živný \(2011\)](#). This paper considers relational constraints only.

where the r -th component is d must equal $\mu_{y,d}$. The remaining constraints ensure that μ and λ define valid probability distributions.

If the variables were constrained to be integers (0 or 1), the solution would correspond exactly to the maximum weight assignment. Without integrality constraints, the BLP provides an upper bound and an approximate solution.

The first step of the proof relies on the following dichotomy result of [Thapper and Živný \(2016\)](#) for finite-valued constraint languages.

THEOREM 1 ([Thapper and Živný 2016](#)). *Let D be an arbitrary finite set and let Γ be a finite-valued constraint language defined on D . Then $\text{VCSP}(\Gamma)$ is tractable if and only if the BLP solves $\text{VCSP}(\Gamma)$. Otherwise, $\text{VCSP}(\Gamma)$ is NP-hard.*

Theorem 1 connects the BLP to tractability of the optimization problem. The second step that completes the proof is the fact that the BLP itself has a direct connection to the existence of market equilibria, which can be characterized via the dual linear program and complementary slackness. This is a similar argument to the one found in [Bikhchandani and Mamer \(1997\)](#), [Bikhchandani et al. \(2001\)](#) and [Bikhchandani and Ostroy \(2002\)](#) for package auctions. [Blumrosen and Nisan \(2007\)](#) interpret the two directions of this step as variants of the first and second welfare theorems for indivisible goods. Indeed, under quasilinearity, if a market equilibrium exists, then prices can be chosen to support any efficient allocation in equilibrium. The complete proof is in the [Appendix](#).

3.2 Illustrative examples

In the following examples, markets are illustrated using graphs G where vertices represent goods Y and edges $E(G)$ represent pairwise relationships (scopes of size 2). The weights of the edges encode the utility gain from a specific assignment of the pair.

First, consider an example with complementary goods called the cutlery auction ([Brandenburg et al. \(2023\)](#)).

EXAMPLE 1 (Cutlery auction). A market has 3 goods $Y = \{S, K, F\}$ (Spoon, Knife, and Fork) and 3 buyers $B = \{\text{Steak}, \text{Fruit}, \text{Spaghetti}\}$. The domain is $D = B$ (assignment of goods to buyers). The buyers have complementary preferences:

- **Steak** values only the pair (F, K) (Fork and Knife).
- **Fruit** values only the pair (K, S) (Knife and Spoon).
- **Spaghetti** values only the pair (S, F) (Spoon and Fork).

Each buyer derives a utility of 1 if they receive exactly their favorite pair (it is assigned to them), and 0 otherwise. $\diamond$

This market is illustrated in Figure 2. The vertices are goods, and edges represent the pairwise complementary scopes. There is a gain in total utility of 1 if both goods at the vertices are assigned to the buyer with the corresponding favorite combination.




FIGURE 2. Cutlery auction

This type of preferences can be captured in the market classes framework by a function that encodes complementarity. Let $\phi^{C,i} \in \mathcal{F}_2$ be defined as:

$$\phi^{C,i}(a_1, a_2) = \begin{cases} C & \text{if } a_1 = a_2 = i, \\ 0 & \text{otherwise,} \end{cases}$$

for $C \in \mathbb{Q}_+$ and $i \in B$. Let $\phi = \{\phi^{C,i} \mid C \in \mathbb{Q}_+, i \in B\}$. The market in Example 1 belongs to $\text{Market}(\phi)$. Specifically, the Steak buyer has one component $((F, K), \phi^{1, \text{Steak}})$, and similarly for the others.

The class $\text{Market}(\phi)$ is the class of all markets where players derive value from obtaining pairs of goods. Each such market is given by a graph G and components $((x(1), x(2)), \phi^{C,i})$ for every edge $(x(1), x(2)) \in E(G)$, buyer i and utility gain C .

The specific instance in Example 1 and Figure 2 has one component $((x_b), \phi^{1,b})$ for every buyer, where x_b is the corresponding edge. This market has no market equilibrium. Any efficient allocation must assign exactly one pair to the corresponding buyer (total utility 1). However, the BLP relaxation for this instance admits a fractional solution with objective value 1.5 (by setting $\mu_{y,b} = 0.5$ for the relevant goods and buyers). Since the BLP overestimates the optimal integer value, the problem is not solved by the BLP. By Theorem 1, the class $\text{VCSP}(\phi)$ is intractable. By Proposition 1, this implies that markets in $\text{Market}(\phi)$ do not robustly possess equilibria.

The failure of equilibrium existence can be seen directly via the dual linear program. Let $p(y, d, b)$ denote the price buyer b pays if good y is assigned to d (here $D = B$, so d is a buyer) as in the definition of the price system. Let u_b denote the payoff variable for buyer b 's component (dual to the primal constraint $\sum \lambda = 1$), and v_y denote the seller's payoff (dual to $\sum \mu = 1$).

The dual constraints corresponding to the primal variables are:

$$\min \sum_{b \in B} u_b + \sum_{y \in Y} v_y$$

s.t. for all $d \in D$,

$$v_K \geq p(K, d, \text{Steak}) + p(K, d, \text{Fruit}),$$

$$v_S \geq p(S, d, \text{Spaghetti}) + p(S, d, \text{Fruit}),$$

$$v_F \geq p(F, d, \text{Steak}) + p(F, d, \text{Spaghetti}),$$

and for all $d_1, d_2 \in D$,

$$u_{\text{Steak}} + p(K, d_1, \text{Steak}) + p(F, d_2, \text{Steak}) \geq \phi^{1, \text{Steak}}(d_1, d_2),$$

$$u_{\text{Fruit}} + p(K, d_1, \text{Fruit}) + p(S, d_2, \text{Fruit}) \geq \phi^{1, \text{Fruit}}(d_1, d_2),$$

$$u_{\text{Spaghetti}} + p(S, d_1, \text{Spaghetti}) + p(F, d_2, \text{Spaghetti}) \geq \phi^{1, \text{Spaghetti}}(d_1, d_2).$$

The complementary slackness conditions imply:

- (i) For every good $y \in Y$ and domain value $d \in D$, if $\mu_{y,d} > 0$, then the seller's payoff must equal the sum of transfers:

$$v_y = \sum_{b \in B: y \in x_b} p(y, d, b).$$

(Here, the inner sum is over all buyers b whose only component's scope x_b contains good y).

- (ii) For every buyer $b \in B$, component t , and local assignment $\mathbf{d} \in D^{x_b}$, if $\lambda_{b,\mathbf{d}} > 0$, then the buyer's component payoff must equal the valuation minus transfers:

$$u_b = \phi^{1,b}(\mathbf{d}) - \sum_{r=1}^{k_b} p(x_b(r), \mathbf{d}(r), b).$$

(Here, $x_b(r)$ is the r -th good in scope x_b and $\mathbf{d}(r)$ is the r -th value in assignment $\mathbf{d}$).

The dual program together with the complementary slackness encodes the market equilibrium conditions in Definition 2. The gap between the BLP value (1.5) and the integer optimum (1) proves that no prices can support the efficient allocation and correspond to a well-defined allocation (an integer solution) for the primal at the same time. This connection between the dual and the market equilibrium is what gives rise to Proposition 1. It allows one to say that the market equilibrium is not guaranteed to exist without looking for the counterexample or analyzing the linear programs. Since $\text{Market}(\phi)$ includes the 'Multiway Cut'⁶ problem, which is known to be NP-hard, Proposition 1 already says that market equilibrium will not exist for some markets in $\text{Market}(\phi)$.

The second example illustrates substitutable and complementary public goods.

EXAMPLE 2 (Public goods with externalities). Consider the problem of allocating a subset of public goods Y . The domain is $D = \{0, 1\}$, where 1 means the good

⁶As they are not directly relevant, I omit the definitions of the standard optimization problems such as 'Minimum Cut', 'Multiway Cut', etc. Their exact statements can be found in complexity theory textbooks or in the VCSP literature (Kolmogorov et al. (2015)).

is provided. Buyers derive utility (or disutility) from individual goods and from pairs of goods (externalities). $\diamond$

Pairwise interactions are once again represented by an undirected graph $G = (Y, E)$. For each edge $(x(1), x(2)) \in E$, include a component with scope $(x(1), x(2))$ and a subutility function $\eta^C \in \mathcal{F}_2$ defined as:

$$\eta^C(a_1, a_2) = \begin{cases} C & \text{if } a_1 = a_2 = 1, \\ 0 & \text{otherwise,} \end{cases}$$

for some $C \in \mathbb{Q}$. Let $\eta_+ = \{\eta^C \mid C > 0\}$ (positive externalities) and $\eta_- = \{\eta^C \mid C < 0\}$ (negative externalities).

Individual good values are captured by unary functions $\psi^{C,1} \in \mathcal{F}_1$:

$$\psi^{C,1}(a) = \begin{cases} C & \text{if } a = 1, \\ 0 & \text{otherwise.} \end{cases}$$

Let $\psi = \{\psi^{C,1} \mid C \in \mathbb{Q}\}$, $\psi_+ = \{\psi^{C,1} \mid C > 0\}$, and $\psi_- = \{\psi^{C,1} \mid C < 0\}$. Sellers have costs for providing goods. If the cost of providing good y is C_y , the seller's payoff is $V_y(1) = -C_y$ or $V_y(0) = 0$. All such payoffs are captured by ψ_- .

Negative Externalities. If the market contains negative externalities ($\text{Market}(\eta_-, \psi)$), equilibria may not exist. The class $\text{Market}(\eta_-, \psi)$ includes the problem of finding a maximum independent set, the largest set of vertices that are not connected to each other, which is NP-hard. By Proposition 1, there exist markets in this class with no equilibrium.

Positive Externalities. If the market contains only positive externalities ($\text{Market}(\eta_+, \psi)$), then the corresponding optimization problem is a supermodular Quadratic Pseudo-Boolean Optimization (QPBO) problem. Minimizing a quadratic submodular function is equivalent to finding a minimum cut in a graph, which is solvable in polynomial time (Ivănescu (1965), Živný (2009)). Therefore, by Proposition 1, every market with only positive pairwise externalities admits a market equilibrium. The equilibrium in this market is similar to Lindahl equilibrium,

where public goods are financed through personalized prices. An example of such a problem is joint financing of a technology with network effects, for example electric car chargers. Installation of individual chargers is costly, but may have positive local and network effects.

EXAMPLE 3 (Graphical complements). Consider again the allocation of public goods Y . Each good has an associated cost and provides some separable utility. Additionally, there are prerequisite constraints: some goods require others to be provided. These prerequisites are represented by a directed graph $G = (Y, E)$, representing a technology tree, where an edge $(x(1), x(2))$ implies that $x(1) \in Y$ is a prerequisite for $x(2) \in Y$. Failure to provide the prerequisite $x(1)$ while providing $x(2)$ incurs a penalty (or is impossible). $\diamond$

This problem is modeled using the unary functions ψ (defined in Example 2) for the separable part, plus an additional set of functions capturing the prerequisite penalties. For each edge $(x(1), x(2)) \in E$, include a component with scope $(x(1), x(2))$ and a penalty function $\zeta^{-C} \in \mathcal{F}_2$ defined as:

$$\zeta^{-C}(a_1, a_2) = \begin{cases} -C & \text{if } a_1 = 0 \text{ and } a_2 = 1, \\ 0 & \text{otherwise,} \end{cases}$$

for some penalty $C \in \mathbb{Q}_+$. Let $\zeta_- = \{\zeta^{-C} \mid C \in \mathbb{Q}_+\}$. If C is sufficiently large, this effectively enforces the logical implication $x(2) = 1 \implies x(1) = 1$.

The optimization problem for the class $\text{Market}(\zeta_-, \psi)$ corresponds to the Maximum Weight Closure problem (or equivalently, Minimum Cut with terminals). This problem is solvable in polynomial time by reduction to a maximum flow problem (Picard (1976)). Therefore, by Proposition 1, every market with graphical prerequisites admits a market equilibrium.

4. CONSERVATIVE MARKETS

In many market design settings, it is natural to assume that while interactions between goods (substitutability, complementarity) may follow a specific structure, the valuations for individual goods can be arbitrary. For example, a designer

might know that goods are pairwise complements, but the specific base values of the goods vary across buyers and instances. Such classes of problems are conservative market classes.

DEFINITION 6 (Conservative market class). A market class $\text{Market}(\mathcal{U})$ is *conservative* if $\mathcal{U}$ contains all unary functions, i.e., $\mathcal{F}_1 \subseteq \mathcal{U}$.

Similarly, a language Γ is called conservative if it contains all unary functions $\mathcal{F}_1$.

For conservative classes, the boundary between existence and non-existence of equilibria is sharply characterized by supermodularity. A function $f : D^k \rightarrow \mathbb{Q}$ is supermodular on a lattice $\mathcal{L} = (D, \leq)$ if for all $\mathbf{d}_1, \mathbf{d}_2 \in D^k$:

$$f(\mathbf{d}_1 \vee \mathbf{d}_2) + f(\mathbf{d}_1 \wedge \mathbf{d}_2) \geq f(\mathbf{d}_1) + f(\mathbf{d}_2),$$

where $\vee$ and $\wedge$ denote the coordinate-wise join (max) and meet (min) under the ordering $\leq$. A market class is supermodular if there exists a total ordering on the domain D such that every function in $\mathcal{U}$ is supermodular with respect to that ordering.

Maximization of supermodular (minimization of submodular) functions is known to be tractable (Iwata et al. (2001)). The following theorem (Theorem 3.6 in Kolmogorov and Živný (2013)) states that for conservative finite-valued languages, these are the only tractable functions:

THEOREM 2 (Kolmogorov and Živný (2013)). *If a conservative finite-valued language is supermodular on some chain, then it is tractable. Otherwise it is NP-hard.*

Combining Proposition 1 with Theorem 2 gives the following characterization for conservative markets.

PROPOSITION 2. *Let $\text{Market}(\mathcal{U})$ be a conservative market class (i.e., $\mathcal{F}_1 \subseteq \mathcal{U}$). Then every market in $\text{Market}(\mathcal{U})$ admits a market equilibrium if and only if there exists a total ordering on D such that every function in $\mathcal{U}$ is supermodular with respect to this ordering.*

Proposition 2 is a strong negative statement. Excluding arbitrary unary terms is usually difficult: buyers may differ substantially in their standalone values for individual goods. Thus, conservativeness leaves very little room for positive equilibrium-existence results beyond additivity.

5. EXTENSION: STRUCTURAL MARKETS

The previous section established that the algebraicity assumption (Remark 2) often leads to non-existence results: while conservativeness seems like an innocuous assumption, it already leads to a very restrictive supermodularity requirement in Proposition 2. However, in many applications, the structure of utility is restricted: for instance, all buyers may agree on which goods are complements. To capture this, the following definition relaxes the requirement that each buyer's component must strictly belong to the class $\mathcal{U}$, and instead requires that the *aggregate* valuation functions belong to $\mathcal{U}$, subject to additional restrictions in a private-goods setting.

DEFINITION 7 (Aligned private-goods market class). Let $D = B$. A market $M = \langle B, Y, D, (U_b)_{b \in B}, (V_y)_{y \in Y} \rangle$ belongs to the aligned private-goods market class $\text{Market}^{\text{al}}(\mathcal{U})$ if

- (i) $V_y \in \mathcal{U}$ for every $y \in Y$;
- (ii) there exist aggregate components $\{(x_t, F_t)\}_{t \in T}$, where $F_t \in \mathcal{U}$, and $x_t = (x_t(1), \dots, x_t(k_t))$, sets $T_b \subseteq T$ and functions $f_{b,t} : D^{k_t} \rightarrow \mathbb{Q}$ for $b \in B$, $t \in T_b$, such that:

- (a) **Decomposition.** For every buyer $b \in B$ and every allocation $\omega \in \Omega$,

$$U_b(\omega) = \sum_{t \in T_b} f_{b,t}(\omega(x_t));$$

- (b) **Aggregation.** For every component $t \in T$ and every $\mathbf{d} \in D^{k_t}$,

$$F_t(\mathbf{d}) = \sum_{\substack{b \in B: \\ t \in T_b}} f_{b,t}(\mathbf{d});$$

(c) **Private goods.** For every buyer $b \in B$, every component $t \in T_b$, and every local assignment $\mathbf{d} \in D^{k_t}$,

$$f_{b,t}(\mathbf{d}) = \begin{cases} C_{b,t}, & \text{if } \mathbf{d} = (b, \dots, b), \\ 0, & \text{otherwise,} \end{cases} \quad \text{for some } C_{b,t} \geq 0.$$

Condition (a) mirrors the one in the definition of a market. Condition (b) allows functions in the components to correspond to sums of utilities of several buyers. Condition (c) ensures that while the aggregate function F_t defines the potential welfare of several buyers, the buyers themselves only care about whether they get the goods in the scope or not.

DEFINITION 8 (Optimization problem for an aligned market class). The optimization problem for a market in $\text{Market}^{\text{al}}(\mathcal{U})$ is to find an allocation $\omega \in \Omega$ that maximizes the objective function $W : \Omega \rightarrow \mathbb{Q}$ defined by:

$$W(\omega) = \sum_{b \in B} U_b(\omega) + \sum_{y \in Y} V_y(\omega(y)) = \sum_{t \in T} F_t(\omega(x_t)) + \sum_{y \in Y} V_y(\omega(y)).$$

Therefore complexity, and ultimately existence of equilibria, is determined by welfare maximization in terms of the aggregate functions F_t . Similarly to market classes, a structural market class $\text{Market}^{\text{al}}(\mathcal{U})$ is said to be tractable if for any finite $\mathcal{U}' \subseteq \mathcal{U}$ there is an algorithm that computes an optimal allocation for any market in $\text{Market}^{\text{al}}(\mathcal{U}')$ in time polynomial in the number of goods $|Y|$.

Aligned market classes prove to have the same existence condition as regular market classes, which is linked to computational tractability in the same way as in Proposition 1.

PROPOSITION 3. *Every market in the structural market class $\text{Market}^{\text{al}}(\mathcal{U})$ admits a market equilibrium if and only if the corresponding optimization problem can be solved in time polynomial in the number of goods.*

5.1 Example: Complements for all buyers

Example 1 featured extreme heterogeneity of preferences: what some considered complements, others considered worthless. Example 4 below shows that even

if buyers agree on the complementarity structure, equilibrium existence is not guaranteed. However, it always exists for such a market with only two buyers.

EXAMPLE 4 (Broadcasters and Subscribers). Consider a market where $|B|$ buyers (Broadcasters) compete to serve $|Y|$ sellers (Subscribers). The subscribers are arranged on a graph $G = (Y, E)$. Broadcasters have costs to serve each subscriber, but can economize if they serve adjacent subscribers (nodes connected by an edge). Specifically, if a broadcaster b serves both $x(1)$ and $x(2)$ where $(x(1), x(2)) \in E$, they gain a cost reduction of $C > 0$. $\diamond$

This fits the structural class framework. The domain is once again $D = B$. First, define a unary function $\psi^{C,i} \in \mathcal{F}_1$ as:

$$\psi^{C,i}(d) = \begin{cases} C & \text{if } d = i, \\ 0 & \text{otherwise.} \end{cases}$$

Let $\psi_{|B|} = \{\psi^{C,i} \mid C \in \mathbb{Q}_-, i \in B\}$ denote the set of negative cost functions (broadcasters' costs to provide service to individual consumers).

For each edge $(x(1), x(2)) \in E$, define a component with scope $(x(1), x(2))$ and aggregate function $\sigma^C \in \mathcal{F}_2$ defined as:

$$\sigma^C(a_1, a_2) = \begin{cases} C & \text{if } a_1 = a_2, \\ 0 & \text{otherwise.} \end{cases}$$

Broadcasters' utilities are composed by functions $f_{b,t}$ for every $b \in B$ and edge $t = (x(1), x(2)) \in E$:

$$f_{b,t}(a_1, a_2) = \begin{cases} C & \text{if } a_1 = a_2 = b, \\ 0 & \text{otherwise.} \end{cases}$$

This satisfies Aggregation ($\sum_b f_{b,t} = \sigma^C$) and Exclusivity (only b gets the value if the condition holds).

The classical problem of finding a minimum cut in a graph is a special case of this problem. This class of problems is known in theoretical computer science

as ‘Potts model with external field’ ([Gridchyn and Kolmogorov \(2013\)](#)), originally used in statistical mechanics to explain ferromagnetism.

- If $|B| = 2$, the problem is the ‘Ising Model’ (equivalent to Minimum Cut), which is tractable. By Proposition 3, a market equilibrium always exists.
- If $|B| \geq 3$, the problem encodes the ‘Multiway Cut’ problem, which is NP-hard. Thus, by Proposition 3, there exist markets in this class (with $|B| \geq 3$) that possess no market equilibrium.

Figure 3 illustrates one market for $|B| = 3$ without an equilibrium. All weights are either 1 (if not marked) or 4 (if marked). Further, the three vertices labelled y_1, y_2, y_3 are particularly desired, respectively, by players 1, 2 and 3 for a utility of 1. The other goods only create value in combination with other goods. The overall problem consists of $(y_i, \psi^{1,i})$ for $i \in \{1, 2, 3\}$ and $((x(1), x(2)), \sigma^1)$ or $((x(1), x(2)), \sigma^4)$ for every $(x(1), x(2)) \in E(G)$.⁷




FIGURE 3. A market with complementary preferences and without a market equilibrium

⁷The example is based on [Papadimitriou et al. \(2008\)](#), specifically it is the modified Graph C with an additional edge.

It is trivial, if tedious, to check that in this example the linear program and the integer program solutions also do not coincide. However, there is also a structural reason for this difference between $|B| = 2$ and $|B| > 2$. The functions in the model with $|B| = 2$ are supermodular, while for $|B| > 2$ they are not. Section 4 showed that when ψ is present among the subutility functions, supermodularity is the exact condition for tractability and, by means of Proposition 1, existence of equilibria.

Preferences over broadcasters. One might ask if the complexity changes for subscribers who have heterogeneous preferences over broadcasters. Let ι be the constant unary function $\iota(d) = C$. A positive valuation $\psi^{C,i}$ with $C > 0$ can be expressed as a combination of negative costs and a constant:

$$\psi^{C,i}(d) = \sum_{j \in B \setminus \{i\}} \psi^{-C,j}(d) + \iota(d).$$

Intuitively, giving buyer i a bonus of C is equivalent to penalizing everyone else by C . By the algebraic properties of VCSPs (Theorem 3 in the Appendix), the class $\text{Market}^{\text{al}}(\sigma \cup \{\psi^{C,i} \mid C \in \mathbb{Q}, i \in B\}) = \text{Market}^{\text{al}}(\sigma \cup \psi_{|B|} \cup \{\psi^{C,i} \mid C \in \mathbb{Q}_+, i \in B\})$ has the same complexity as $\text{Market}^{\text{al}}(\sigma \cup \psi_{|B|})$. Unary functions from $\{\psi^{C,i} \mid C \in \mathbb{Q}_+, i \in B\}$ capture (positive) utility that consumers obtain from particular broadcasters. Thus, allowing heterogeneous subscriber preferences does not affect equilibrium existence. Note also that since prices can be negative, it is always possible to obtain a new equilibrium by adding or subtracting a constant from all prices (Remark 1). Then one can obtain an equilibrium in which subscribers pay the broadcasters instead of vice versa.

6. ITERATIVE EXCHANGE

A natural economic mechanism can be used to turn the search for a market equilibrium into an iterative procedure similar to an ascending auction. To describe this procedure, I define an auxiliary linear program called the Restricted Primal Program (RPP). I will assume that each buyer's U_b is represented by exactly one component (x_b, f_b) with $k_b = |x_b| = ar(f_b)$. This assumption does not change the

complexity of the problem.⁸ Let $\hat{\mathcal{D}}_b(p)$ be the demand set of buyer b in terms of allocations of goods relevant to b , i.e. x_b , instead of the whole allocation ω . Formally, $\hat{\mathcal{D}}_b(p) = \{\mathbf{d} : x_b \rightarrow D, \omega(x_b) = \mathbf{d} \text{ for some } \omega \in \mathcal{D}_b\}$. The RPP is then:

$$\begin{aligned}
 & \min \sum_{y \in Y} s_y + \sum_{b \in B} s_b + \sum_{b \in B} \sum_{d \in D} \sum_{r=1}^{k_b} s_{b,r,d}, \\
 & \text{s.t.} \quad \sum_{\substack{\mathbf{d} \in \hat{\mathcal{D}}_b(p), \\ \mathbf{d}(r)=d}} \lambda_{b,\mathbf{d}} + s_{b,r,d} = \mu_{y_b,r,d} \quad \forall b \in B, r \in \{1, \dots, k_b\}, d \in D, \\
 & \quad \sum_{d \in \mathcal{S}_y(p)} \mu_{y,d} + s_y = 1 \quad \forall y \in Y, \\
 & \quad \sum_{\mathbf{d} \in \hat{\mathcal{D}}_b(p)} \lambda_{b,\mathbf{d}} + s_b = 1 \quad \forall b \in B, \\
 & \quad \lambda, \mu, s_y, s_b, s_{b,r,d} \geq 0.
 \end{aligned} \tag{2}$$

Here, $s_y, s_b, s_{b,r,d}$ are slack variables. The RPP is parametrized by sets $\hat{\mathcal{D}}_b$ and $\mathcal{S}_y$. The set $\hat{\mathcal{D}}_b$ represents the demand set of buyer b — the set of allocations that maximizes the payoff of b , or equivalently the set of constraints of the dual corresponding to the optimization by b that bind. Similarly, the set $\mathcal{S}_y$ is the supply set of seller y . A demand set $\hat{\mathcal{D}}_b$ is said to be satisfied if $\lambda_{x_b,d^x} = 1$ for some $\mathbf{d} \in \hat{\mathcal{D}}_b$. A supply set $\mathcal{S}_y$ is satisfied if $\mu_{y,d} = 1$ for some $d \in \mathcal{S}_y$. The RPP represents a search for a feasible solution to LP (1), while satisfying the implications of complementary slackness. If the minimum objective value is 0, the solution is also a solution to the BLP (and thus an equilibrium if it is integer).

The dual of the RPP (2), called the Restricted Dual Program (RDP), searches for a direction of price adjustment that improves the solution to the BLP. The

⁸This auxiliary result is formally stated as Theorem 3 in the [Appendix](#).

variables $\Delta p, \Delta v, \Delta u$ correspond to changes in prices and payoffs:

$$\begin{aligned}
 \max \Delta W &= \sum_{y \in Y} \Delta v_y + \sum_{b \in B} \Delta u_b, \\
 \text{s.t. } \Delta v_y &\leq \sum_{\substack{b \in B: \\ y \in x_b}} \Delta p(y, d, b) \quad \forall y \in Y, d \in \mathcal{S}_y(p), \\
 \Delta u_b &\leq - \sum_{r=1}^{k_b} \Delta p(x_b(r), \mathbf{d}(r), b) \quad \forall b \in B, \mathbf{d} \in \hat{\mathcal{D}}_b(p), \\
 \Delta v_y, \Delta u_b, \Delta p &\leq 1.
 \end{aligned} \tag{3}$$

The RDP identifies a direction Δp that decreases the dual objective function of the BLP. The constraints of the RDP ensure that prices change in the direction where the slack of the constraints permits. For separable preferences, this reduces to the standard logic of ascending auctions: prices increase for over-demanded goods, where constraints are tight. Then the increase in price would improve v_y for each seller, but decrease u_b for all buyers who demand it, making the overall effect on the objective negative.

The iterative procedure is summarized in Algorithm 1. Importantly, it relies only on demand/supply queries: agents report their optimal sets $\hat{\mathcal{D}}_{b,t}(p)$ and $\mathcal{S}_y(p)$, without having to report full utility functions. The algorithm can be seen as a primal–dual algorithm for solving the BLP for the market. This connection yields the following convergence result.

PROPOSITION 4. *The iterative procedure in Algorithm 1 terminates with a market equilibrium if one exists.*

The primal–dual connection also suggests a natural upper bound on the number of iterations. In a standard primal–dual algorithm, one chooses the maximal step size that preserves dual feasibility, so that each iteration makes at least one

Algorithm 1: Iterative Exchange Procedure

Input: Market M , initial prices $p \equiv 0$, step size $\epsilon > 0$.

1 while true do

 // Step 1: Demand/Supply Queries

2 For each buyer b , collect $\hat{D}_b(p)$.

3 For each seller y , collect $S_y(p)$.

 // Step 2: Solve Restricted Dual

4 Solve RDP with current sets to find direction Δp and optimal value ΔW^* .

5 if $\Delta W^* \leq 0$ (No descent direction) then

 // Step 3: Check Primal Feasibility

6 Solve RPP.

7 if objective is 0 and solution is integer then

8 Find allocation ω supported by RPP solution.

9 return Equilibrium (ω, p) .

10 else

11 return "No market equilibrium exists."

12 else

 // Step 4: Update Prices

13 $p \leftarrow p - \epsilon \cdot \Delta p$.

 // (Optional) Normalize prices for all $y \in Y$, $b \in B$,
 and $d \in D$:

14 $p(y, d, b) \leftarrow p(y, d, b) + \epsilon \cdot c(y, b)$ for some $c: Y \times B \rightarrow \mathbb{Q}$.

previously slack dual constraint tight. Since the dual contains $|Y||D|$ seller constraints and $\sum_{b \in B} |D|^{k_b}$ buyer constraints, this gives an upper bound of

$$|Y||D| + \sum_{b \in B} |D|^{k_b}$$

productive iterations. Algorithm 1 does not, however, have access to enough information to implement this step-size rule, because buyers report only demand sets and not their utility values on all bundles. Accordingly, the bound above applies to “productive” iterations—those in which the reported sets change—rather than to the total number of auction rounds. The total number of rounds may be larger, depending on how accurately the designer selects ϵ : it must be small enough to preserve feasibility, but large enough to induce a change in the reported sets.

By Remark 1, the optional renormalization step in the algorithm (Step 4) can be used to constrain the price path. For instance, one can ensure that the price of a specific ‘endowment’ allocation remains zero, or that prices remain non-negative, by shifting all prices by a uniform constant $c_{y,b}$ at each step. This does not affect the incentives or the direction of the subgradient Δp , but allows the procedure to mimic specific auction formats (e.g., ascending prices) or choose the direction of transfers as in Example 4.

7. DISCUSSION

7.1 *Infinite VCSP*

This paper explicitly considers only finite-valued VCSP, which corresponds to an economic assumption of all allocations being feasible. This leaves many problems out of scope, for example capacity constraints. A similar analysis can be done for the infinite-valued VCSP dichotomy, but the equivalence between equilibrium existence and tractability breaks in a subtle way: the intractability may hide in the problem of finding a feasible allocation, not in solving for an efficient one. In finite-valued VCSP, producing a feasible allocation is trivial. Perhaps surprisingly, the general-valued algorithm is simply the ‘feasibility preprocessing followed by BLP’. Thus results here should apply with minimal adjustment to infinite-valued classes where the feasibility problem is known to be tractable.

7.2 OXS and unit-demand

Some positive equilibrium-existence results rely on structural restrictions that are not encoded by $\text{Market}(U)$. Consider two standard examples: unit-demand and OR-of-XORs-of-singletons (OXS), both known to guarantee robust equilibrium existence ([Lehmann et al. \(2006\)](#)).

Let $S_b(\omega) = \{y \in Y : \omega(y) = b\}$ be the set of goods allocated to buyer b . A unit-demand utility function is a function U_b of the form $U_b(\omega) = \max_{y \in x_b \cap S_b(\omega)} C_y$ for some set of goods $x_b \subseteq Y$ and some constants $C_y \in \mathbb{Q}$ for each $y \in Y$. Buyer b with such U_b only extracts utility from a single most-valued good allocated to b . It would seem natural to incorporate this into the algebraic framework by defining a market class with (x_b, f_b) for every buyer, where $f_b = \max_{y \in x_b \cap S_b(\omega)} C_y$. And while all unit-demand markets would be part of this class, by Definition 3, it would also necessarily include markets where valuations are *sums* of unit-demand maximum functions. Unfortunately, this class is intractable and thus equilibria do not exist for some markets. Therefore, the approach is silent about unit-demand markets.

The distinction is also visible in the case of OXS valuations. A buyer's valuation U_b is OXS if there exist k unit-demand valuations with weights $C_y^j \in \mathbb{Q}$ and scopes $x_j, j = 1, \dots, k$, such that

$$U_b(\omega) = \max_{\substack{S^1, \dots, S^k \subseteq S_b(\omega) \\ S^i \cap S^j = \emptyset \text{ for } i \neq j}} \sum_{j=1}^k \max_{y \in S_j \cap x_j} C_y^j.$$

Note the disjointness requirement: the unit-demand clauses may value overlapping sets of goods, but in the optimizing representation each allocated good may be assigned to at most one clause. Thus OXS is not a pointwise sum of unit-demand valuations either. The disjointness requirement is not captured by the market class defined above and cannot be captured under algebraicity. The fact that the class is otherwise intractable shows that this is a crucial requirement for equilibrium existence for OXS.

Both problems have the same explanation: the tractability of these classes crucially rests on restricting the structure of the allocation or utilities. The bipartite nature of such markets cannot be encoded within the algebraic framework.

8. CONCLUSION

This paper explores the three-way connection between complexity theory, linear programming, and market equilibria. The aim is to inform market design by identifying classes of valuations that ensure the existence of equilibria. Proposition 1 is a step in this direction, offering an alternative view on the relationship between tractability and equilibria beyond unimodularity.

Proposition 2 also has an important negative implication for market design. It shows that for markets in which all unary valuations are allowed, a restrictive form of supermodularity is the only structural property that guarantees equilibrium existence. For private goods, this condition requires utilities to be supermodular with respect to some ordering of the buyers, which is highly restrictive. Thus, if additively separable components cannot be excluded, there is very limited scope for general equilibrium-existence guarantees beyond simple additivity, and one must look beyond global properties such as conservativeness.

Moreover, several classes of valuations known to admit equilibria do not align with the algebraic structures considered here. Foremost, assignment markets as studied by [Shapley and Shubik \(1971\)](#) rely on the bipartite nature of interactions, which require structural constraints absent from the relational VCSP framework. Exploring this ‘other side’ ([Carbonnel et al. \(2022\)](#)) of valued constraint satisfaction problems could yield further positive results for new classes of valuations.

APPENDIX: REMAINING PROOFS

The following definitions and results will be used in the proofs. These are adapted from the VCSP literature ([Cohen and Jeavons \(2006\)](#), [Živný \(2009\)](#)). For any VCSP instance $I = \langle Y, D, f \rangle$ and any ordered tuple of distinct variables $x = (x(1), \dots, x_k) \in Y^k$, the *projection* of I onto x , denoted f^x , is defined as:

$$f^x(\mathbf{d}) = \max\{f(\omega) \mid \omega : Y \rightarrow D \text{ such that } \omega(x) = \mathbf{d}\}.$$

Intuitively, a projection is a restriction of the cost function of the instance to a smaller subset of variables. A function $\theta : D^k \rightarrow \mathbb{Q}$ is *expressible* over a valued constraint language Γ if there exists an instance $I \in \text{VCSP}(\Gamma)$ and a tuple of variables x in I such that $\theta = f^x$.

The following result from Živný (2009), which is a tightening of the original polynomial-time equivalence in Cohen and Jeavons (2006), establishes that adding functions to a language Γ that are already expressible in it does not significantly change its complexity.

THEOREM 3 (Živný 2009, Theorem 2.3.4). *For any valued constraint language Γ and any function θ expressible over Γ , the problems $\text{VCSP}(\Gamma)$ and $\text{VCSP}(\Gamma \cup \{\theta\})$ are log-space equivalent.*

This invariance translates to a similar notion for the economic setting. If a certain welfare function can be constructed by combining the utilities of existing agents (e.g., a "coalition" maximizing their joint payoff), introducing an agent with explicitly that preference structure does not alter the existence of equilibrium.

Additionally, the computational complexity of a language is invariant under positive affine transformations (Cohen and Jeavons (2006), Živný (2009)). Adding a constant to a function is a monotone transformation that does not affect the set of optimal solutions. Scaling a function by a positive rational factor is equivalent to having multiple copies of the component, which is covered by Theorem 3. Thus, any language is closed under these operations without loss of generality. Formally, let ι be the constant unary function defined by $\iota(d) = 1$ for all $d \in D$. Then the following is immediate:

REMARK 3. Let Γ be a finite-valued constraint language. Let Γ' be obtained by scaling functions in Γ by rational factors. Then $\text{VCSP}(\Gamma)$ is NP-hard if and only if $\text{VCSP}(\Gamma')$ is NP-hard, and if and only if $\text{VCSP}(\Gamma \cup \{\iota\})$ is NP-hard.

I now prove the main result.

PROOF OF PROPOSITION 1. Consider any market $M = \langle B, Y, D, (U_b)_{b \in B}, (V_y)_{y \in Y} \rangle$ in the class. For each buyer b , let x_b be a scope constructed by ordering the set of all distinct goods that appear in any of buyer b 's components $\{(x_{b,t}, f_{b,t})\}_t$. Let f_b be defined by $f_b(\omega(x_b)) = U_b(\omega)$ be the buyer's utility expressed as a single function on this combined scope. Let $k_b = |x_b| = ar(f_b)$.

Construct a VCSP instance $\langle Y, D, f \rangle$ with objective function

$$f(\omega) = \sum_{b \in B} f_b(\omega(x_b)) + \sum_{y \in Y} V_y(\omega(y)).$$

This instance belongs to $\text{VCSP}(\mathcal{U} \cup \{f_b\}_{b \in B})$. Note that f_b is expressible over $\mathcal{U}$ because it is a sum of functions in $\mathcal{U}$. By Theorem 3, adding $\{f_b\}_{b \in B}$ to the language does not change its tractability. Thus, considering the extended language $\mathcal{U} \cup \{f_b\}_{b \in B}$ is without loss of generality.

The BLP for this instance after simplifying is:

$$\begin{aligned} \max \quad & \sum_{b \in B} \sum_{\mathbf{d} \in D^{x_b}} f_b(\mathbf{d}) \lambda_{b,\mathbf{d}} + \sum_{y \in Y} \sum_{d \in D} V_y(d) \mu_{y,d}, \\ \text{s.t.} \quad & \sum_{\substack{\mathbf{d} \in D^{x_b}: \\ \mathbf{d}(r)=d}} \lambda_{b,\mathbf{d}} = \mu_{x_b(r),d} & \forall b \in B, r \in \{1, \dots, k_b\}, d \in D, \\ & \sum_{d \in D} \mu_{y,d} = 1 & \forall y \in Y, \\ & \sum_{\mathbf{d} \in D^{x_b}} \lambda_{b,\mathbf{d}} = 1 & \forall b \in B, \\ & \lambda, \mu \geq 0. \end{aligned} \tag{4}$$

The dual program of (4) is:

$$\begin{aligned} \min \quad & \sum_{b \in B} u_b + \sum_{y \in Y} v_y, \\ \text{s.t.} \quad & v_y \geq V_y(d) + \sum_{\substack{b \in B: \\ y \in x_b}} p(y, d, b) & \forall y \in Y, d \in D, \\ & u_b \geq f_b(\mathbf{d}) - \sum_{r=1}^{k_b} p(x_b(r), \mathbf{d}(r), b) & \forall b \in B, \mathbf{d} \in D^{x_b}. \end{aligned} \tag{5}$$

Since the BLP solves the instance, there exists an integer primal solution. Let $\omega \in \Omega$ be an allocation, such that $\omega(y) = d$ iff $\mu_{y,d} = 1$. The constructed ω is a well-defined allocation by the constraints of the primal problem. By strong duality, there exists a dual solution (v, u, p) with the same objective value. The complementary slackness conditions imply

$$\mu_{y,d} = 1 \implies v_y = V_y(d) + \sum_{\substack{b \in B: \\ y \in x_b}} p(y, d, b), \quad (6)$$

$$\lambda_{b,d} = 1 \implies u_b = f_b(d) - \sum_{r=1}^{k_b} p(x_b(r), d(r), b). \quad (7)$$

These are exactly the Seller and Buyer optimality conditions for the market equilibrium (Definition 2). Optimization by sellers is equation (6) combined with the first dual constraint, and optimization by buyers is equation (7) combined with the second dual constraint. Thus, (ω, p) is a market equilibrium.

To prove the converse, suppose $\text{Market}(\mathcal{U})$ is intractable. By Theorem 1, since the class is intractable, there exists an instance that is not solved by the BLP, i.e. for which the BLP is not exact. The corresponding market can be represented by a set of components $\{(x_{b,t}, f_{b,t})\}$. Construct a market M where each such component corresponds to one buyer b' with utility $U_{b'}(\omega) = f_{b,t}(\omega(x_{b,t}))$. This market is also in $\text{Market}(\mathcal{U})$. Linear program (4) is the BLP for this market as well since $W(\cdot)$ for the two problems and the respective optimization problems are equivalent. Since the BLP is not exact, the optimal value is strictly larger than the value of every integral feasible solution. If this market admitted a market equilibrium, the equilibrium allocation and prices would generate an integral primal solution and a dual solution satisfying complementary slackness, hence an integral optimum of the BLP. This contradicts non-exactness of the BLP. Therefore this market has no market equilibrium.

□

PROOF OF PROPOSITION 4. Fix an iteration of Algorithm 1 with current prices p .
For each buyer $b \in B$, let

$$u_b = \max_{\mathbf{d} \in D^{k_b}} \left(f_b(\mathbf{d}) - \sum_{r=1}^{k_b} p(x_b(r), \mathbf{d}(r), b) \right),$$

and let $\hat{\mathcal{D}}_b(p)$ be the set of maximizers. For each good $y \in Y$, let

$$v_y = \max_{d \in D} \left(V_y(d) + \sum_{b \in B} p(y, d, b) \right),$$

and let $\mathcal{S}_y(p)$ be the set of maximizers. Then (p, u, v) is feasible for the dual of the BLP (5), and the constraints indexed by $\hat{\mathcal{D}}_b(p)$ and $\mathcal{S}_y(p)$ are binding at (p, u, v) .

Let $(\Delta p, \Delta u, \Delta v)$ be an optimal solution to the RDP ((3)) with objective $\Delta W = \sum_{y \in Y} \Delta v_y + \sum_{b \in B} \Delta u_b$. Because the RDP is the dual of the RPP, strong duality implies $\Delta W \geq 0$. If $\Delta W > 0$, consider the update

$$p^+ = p - \epsilon \Delta p, \quad u^+ = u - \epsilon \Delta u, \quad v^+ = v - \epsilon \Delta v.$$

For sufficiently small ϵ , all dual constraints that were slack at (p, u, v) remain slack at (p^+, u^+, v^+) , while the RDP inequalities ensure that the binding constraints indexed by $\hat{\mathcal{D}}_b(p)$ and $\mathcal{S}_y(p)$ remain satisfied; hence (p^+, u^+, v^+) is dual-feasible. Moreover,

$$\sum_{y \in Y} v_y^+ + \sum_{b \in B} u_b^+ = \sum_{y \in Y} v_y + \sum_{b \in B} u_b - \epsilon \Delta W < \sum_{y \in Y} v_y + \sum_{b \in B} u_b,$$

so the dual objective strictly decreases whenever a price update is performed. Since the demand and supply correspondences are subsets of finite sets, only finitely many RDPs appear. Hence, whenever the RDP optimum is positive, it is bounded away from zero by some common constant $\delta > 0$. Thus each price update decreases the dual objective by at least $\epsilon \delta$. Since the dual objective is bounded below by the BLP optimum, only finitely many updates are possible, and the algorithm eventually terminates.

At termination, the RDP optimum satisfies $\Delta W = 0$, hence the RPP optimum is also 0. Let (λ, μ) be an optimal RPP solution. By construction of the RPP,

(λ, μ) is feasible for the BLP (4) and satisfies the complementary-slackness conditions. Thus (λ, μ) is an optimal BLP solution. If (λ, μ) is integral, define ω by $\omega(y) = d \iff \mu_{y,d} = 1$. Then the dual (5) and complementary slackness conditions yield buyer and seller optimality, so (ω, p) is a market equilibrium. If (λ, μ) is not integral, then the BLP does not solve this instance. By Theorem 1 and Proposition 1, a market equilibrium does not exist. $\square$

I now prove the characterization for the aligned market classes introduced in Section 5.

PROOF OF PROPOSITION 3. Consider a market

$$M = \langle B, Y, D, (U_b)_{b \in B}, (V_y)_{y \in Y} \rangle \in \text{Market}^{\text{al}}(\mathcal{U}).$$

Let $D = B$ and let aggregate components $\{(x_t, F_t)\}_{t \in T}$, functions $f_{b,t}$ and constants $C_{b,t}$ for $b \in B, t \in T_b$ satisfy Definition 7.

Construct the aggregate VCSP instance

$$f(\omega) = \sum_{t \in T} F_t(\omega(x_t)) + \sum_{y \in Y} V_y(\omega(y)).$$

This instance belongs to $\text{VCSP}(\mathcal{U})$.

Now construct another linear program that represents the buyers' utilities separately. It is referred to as a separable linear program (SLP) below.

$$\begin{aligned} \max \quad & \sum_{b \in B} \sum_{t \in T_b} \sum_{\mathbf{d} \in D^{k_t}} f_{b,t}(\mathbf{d}) \lambda_{b,t,\mathbf{d}} + \sum_{y \in Y} \sum_{d \in D} V_y(d) \mu_{y,d} \\ \text{s.t.} \quad & \sum_{\substack{\mathbf{d} \in D^{k_t}: \\ \mathbf{d}(r)=d}} \lambda_{b,t,\mathbf{d}} = \mu_{x_t(r),d} & \forall b \in B, t \in T_b, r = 1, \dots, k_t, d \in D, \\ & \sum_{d \in D} \mu_{y,d} = 1 & \forall y \in Y, \\ & \sum_{\mathbf{d} \in D^{k_t}} \lambda_{b,t,\mathbf{d}} = 1 & \forall b \in B, t \in T_b, \\ & \lambda, \mu \geq 0. \end{aligned}$$

(8)

I first show that BLP and SLP have the same optimal value. Given any feasible solution (λ_t, μ) of the original BLP, define

$$\lambda_{b,t,d} = \lambda_{t,d} \quad \forall b \in B, t \in T_b, d \in D^{k_t}.$$

This gives a feasible solution of SLP (8), and the objective values are equal by aggregation. Hence the value of SLP is at least the value of the BLP.

Conversely, fix any feasible solution (λ, μ) of SLP. For each $t \in T$, define

$$m_{b,t} = \min_{r=1,\dots,k_t} \mu_{x_t(r),b}, \quad R_t = 1 - \sum_{b \in B} m_{b,t}.$$

Then $\mu_{x_t(r),d} - m_{d,t} \geq 0$ and $\sum_{d \in D} (\mu_{x_t(r),d} - m_{d,t}) = R_t$ for all r . I will say that d is constant if it is of the form $(b, \dots, b)$ for some $b \in B$. Define $\hat{\lambda}_t$ by

$$\hat{\lambda}_{t,d} := \begin{cases} m_{b,t}, & \text{if } R_t = 0 \text{ and } d \text{ is constant,} \\ 0, & \text{if } R_t = 0 \text{ and } d \text{ is not constant,} \\ m_{b,t} + R_t \prod_{r=1}^{k_t} \frac{\mu_{x_t(r),d(r)} - m_{d(r),t}}{R_t}, & \text{if } R_t > 0 \text{ and } d \text{ is constant,} \\ R_t \prod_{r=1}^{k_t} \frac{\mu_{x_t(r),d(r)} - m_{d(r),t}}{R_t}, & \text{if } R_t > 0 \text{ and } d \text{ is not constant.} \end{cases}$$

The constructed $\hat{\lambda}_t$ places as much mass as possible on assignments of the form $(b, \dots, b)$ and then distributes the rest to satisfy the constraints. Then $\hat{\lambda}_t \geq 0$, $\sum_{d \in D^{k_t}} \hat{\lambda}_{t,d} = 1$, and, for every $r = 1, \dots, k_t$ and $d \in D$,

$$\sum_{\substack{d \in D^{k_t}: \\ d(r)=d}} \hat{\lambda}_{t,d} = \mu_{x_t(r),d}.$$

For each $b \in B$ with $t \in T_b$, feasibility of SLP gives

$$\lambda_{b,t,(b,\dots,b)} \leq \sum_{\substack{d \in D^{k_t}: \\ d(r)=b}} \lambda_{b,t,d} = \mu_{x_t(r),b} \quad \forall r = 1, \dots, k_t.$$

Hence

$$\lambda_{b,t,(b,\dots,b)} \leq m_{b,t}.$$

Therefore, since $C_{b,t} \geq 0$,

$$\begin{aligned} \sum_{\mathbf{d} \in D^{k_t}} F_t(\mathbf{d}) \hat{\lambda}_{t,\mathbf{d}} &\geq \sum_{\substack{b \in B: \\ t \in T_b}} C_{b,t} m_{b,t} \\ &\geq \sum_{\substack{b \in B: \\ t \in T_b}} C_{b,t} \lambda_{b,t,(b,\dots,b)} \\ &= \sum_{\substack{b \in B: \\ t \in T_b}} \sum_{\mathbf{d} \in D^{k_t}} f_{b,t}(\mathbf{d}) \lambda_{b,t,\mathbf{d}}. \end{aligned}$$

Summing over $t \in T$, and adding the common seller term

$$\sum_{y \in Y} \sum_{d \in D} V_y(d) \mu_{y,d},$$

shows that the BLP has value at least as large as the SLP. Hence the BLP and the SLP have the same optimal value.

By assumption, the BLP solves this instance with an integer solution corresponding to some allocation ω . Let (λ, μ) be the variables that correspond to this solution. This solution is feasible for SLP and corresponds to the same allocation ω . Its objective value is

$$\sum_{b \in B} \sum_{t \in T_b} f_{b,t}(\omega(x_t)) + \sum_{y \in Y} V_y(\omega(y)) = \sum_{t \in T} F_t(\omega(x_t)) + \sum_{y \in Y} V_y(\omega(y)),$$

where the equality follows from the aggregation condition. Hence its value is the optimal value of BLP. Since BLP and SLP have the same optimal value, the induced integral solution is also optimal for SLP.

Conversely, if the aggregate BLP is not exact, then

$$\text{val}(\text{BLP}) > \max_{\omega} W(\omega).$$

Integral feasible solutions of BLP and SLP are induced by the same allocations and have the same objective values. Therefore the best integral value of SLP is

also $\max_{\omega} W(\omega)$. Since

$$\text{val}(\text{SLP}) = \text{val}(\text{BLP}),$$

we obtain

$$\text{val}(\text{SLP}) > \max_{\omega} W(\omega).$$

Thus SLP has no integral optimal solution.

It remains to connect SLP to equilibrium. The dual of SLP is

$$\begin{aligned} \min \quad & \sum_{y \in Y} v_y + \sum_{b \in B} \sum_{t \in T_b} u_{b,t}, \\ \text{s.t.} \quad & v_y \geq V_y(d) + \sum_{\substack{b \in B, t \in T_b: \\ y \in x_t}} p(y, d, b, t) \quad \forall y \in Y, d \in D, \\ & u_{b,t} \geq f_{b,t}(\mathbf{d}) - \sum_{r=1}^{k_t} p(x_t(r), \mathbf{d}(r), b, t) \quad \forall b \in B, t \in T_b, \mathbf{d} \in D^{k_t}. \end{aligned} \tag{SD}$$

By the same argument as in the proof of Proposition 1, a market equilibrium exists if and only if SLP has an integral optimum. Indeed, if SLP has an integral optimum induced by an allocation ω , let (v, u, p) be an optimal dual solution satisfying complementary slackness. Define

$$p^*(y, d, b) = \sum_{\substack{t \in T_b: \\ y \in x_t}} p(y, d, b, t).$$

Complementary slackness and the second family of dual constraints imply buyer optimality at p^* . The first family of dual constraints and complementary slackness imply seller optimality. Thus (ω, p^*) is a market equilibrium.

Conversely, suppose that the aggregate VCSP class generated by $\text{Market}^{\text{al}}(\mathcal{U})$ is not solved by the BLP. Then there is an instance

$$\sum_{t \in T} F_t(\omega(x_t)) + \sum_{y \in Y} V_y(\omega(y))$$

arising from some market $M \in \text{Market}^{\text{al}}(\mathcal{U})$ such that its BLP has no integral optimal solution. By the value-equivalence between the BLP and the SLP established above, the corresponding SLP also has no integral optimal solution. Since these

conditions correspond to the conditions for market equilibrium, a market equilibrium exists if and only if SLP has an integral optimum. Then by the same argument as in Proposition 1, there is no pair of an allocation and a price vector that would be a market equilibrium for this market. $\square$

REFERENCES

- Baldwin, Elizabeth and Paul Klemperer (2019), “Understanding preferences: demand types”, and the existence of equilibrium with indivisibilities.” *Econometrica*, 87(3), 867–932. [4, 5]
- Bikhchandani, Sushil, Sven de Vries, James Schummer, and Rakesh V Vohra (2001), “Linear programming and vickrey auctions.” *IMA Volumes in Mathematics and its Applications*, 127, 75–116. [4, 16]
- Bikhchandani, Sushil and John W. Mamer (1997), “Competitive equilibrium in an exchange economy with indivisibilities.” *Journal of Economic Theory*, 74(2), 385–413, URL <https://www.sciencedirect.com/science/article/pii/S0022053196922693>. [16]
- Bikhchandani, Sushil and Joseph M Ostroy (2002), “The package assignment model.” *Journal of Economic Theory*, 107(2), 377–406. [2, 4, 5, 16]
- Blumrosen, Liad and Noam Nisan (2007), “Combinatorial auctions.” In *Algorithmic Game Theory* (Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani, eds.), chapter 11, 267–300, Cambridge University Press. [16]
- Blumrosen, Liad and Noam Nisan (2010), “On the computational power of demand queries.” *SIAM Journal on Computing*, 39(4), 1372–1391. [3]
- Brandenburg, Marie-Charlotte, Christian Haase, and Ngoc Mai Tran (2023), “Competitive equilibrium always exists for combinatorial auctions with graphical pricing schemes.” *La Matematica*, 2(1), 1–30. [2, 5, 16]
- Candogan, Ozan, Asuman Ozdaglar, and Pablo A Parrilo (2015), “Iterative auction design for tree valuations.” *Operations Research*, 63(4), 751–771. [5]

- 1 Carbonnel, Clément, Miguel Romero, and Stanislav Živný (2022), “The complex- 1
2 ity of general-valued constraint satisfaction problems seen from the other side.” 2
3 *SIAM Journal on Computing*, 51(1), 19–69. [33] 3
- 4 Cohen, David and Peter Jeavons (2006), “The complexity of constraint lan- 4
5 guages.” In *Foundations of Artificial Intelligence*, volume 2, 245–280, Elsevier. [33, 5
6 34] 6
- 7 Cooper, Martin C and Stanislav Živný (2011), “Hybrid tractability of valued con- 7
8 straint problems.” *Artificial Intelligence*, 175(9-10), 1555–1569. [15] 8
9 9
- 10 Delacrétaz, David, Simon Loertscher, and Claudio Mezzetti (2025), “Public goods, 10
11 social alternatives, and the lindahl-vcg relationship.” *Journal of Economic Theory*, 11
12 106055. [5] 12
- 13 Facchinei, Francisco and Christian Kanzow (2010), “Generalized nash equilib- 13
14 rium problems.” *Annals of Operations Research*, 175(1), 177–211. [4] 14
15 15
- 16 Gridchyn, Igor and Vladimir Kolmogorov (2013), “Potts model, parametric 16
17 maxflow and k-submodular functions.” In *Proceedings of the IEEE International* 17
18 *Conference on Computer Vision*, 2320–2327. [3, 26] 18
- 19 Gul, Faruk and Ennio Stacchetti (1999), “Walrasian equilibrium with gross sub- 19
20 stitutes.” *Journal of Economic theory*, 87(1), 95–124. [1, 5] 20
21 21
- 22 Hoffman, Alan J (1956), “Integral boundary points of convex polyhedra.” *Linear* 22
23 *inequalities and related systems*, 223–246. [2] 23
- 24 Ivănescu, Peter L (1965), “Some network flow problems solved with pseudo- 24
25 boolean programming.” *Operations Research*, 13(3), 388–399. [20] 25
26 26
- 27 Iwata, Satoru, Lisa Fleischer, and Satoru Fujishige (2001), “A combinatorial 27
28 strongly polynomial algorithm for minimizing submodular functions.” *Journal* 28
29 *of the ACM (JACM)*, 48(4), 761–777. [22] 29
- 30 Kelso Jr, Alexander S and Vincent P Crawford (1982), “Job matching, coalition for- 30
31 mation, and gross substitutes.” *Econometrica: Journal of the Econometric Society*, 31
32 1483–1504. [1, 4] 32

1 Kolmogorov, Vladimir, Johan Thapper, and Stanislav Živný (2015), “The power of 1
2 linear programming for general-valued csps.” *SIAM Journal on Computing*, 44(1), 2
3 1–36. [19] 3

4 Kolmogorov, Vladimir and Stanislav Živný (2013), “The complexity of conserva- 4
5 tive valued csps.” *Journal of the ACM (JACM)*, 60(2), 1–38. [22] 5
6 6

7 Lehmann, Benny, Daniel Lehmann, and Noam Nisan (2006), “Combina- 7
8 torial auctions with decreasing marginal utilities.” *Games and Economic* 8
9 *Behavior*, 55(2), 270–296, URL [https://ideas.repec.org/a/eee/gamebe/](https://ideas.repec.org/a/eee/gamebe/v55y2006i2p270-296.html) 9
10 [v55y2006i2p270-296.html](https://ideas.repec.org/a/eee/gamebe/v55y2006i2p270-296.html). [5, 32] 10

11 Mishra, Debasis and David C Parkes (2007), “Ascending price vickrey auctions for 11
12 general valuations.” *Journal of Economic Theory*, 132(1), 335–366. [3] 12
13 13

14 Nisan, Noam (2000), “Bidding and allocation in combinatorial auctions.” In *Pro-* 14
15 *ceedings of the 2nd ACM Conference on Electronic Commerce*, 1–12. [5] 15
16 16

17 Nisan, Noam and Ilya Segal (2006), “The communication requirements of ef- 17
18 ficient allocations and supporting prices.” *Journal of Economic Theory*, 129(1), 18
19 192–224. [5] 19

20 Papadimitriou, Christos, Michael Schapira, and Yaron Singer (2008), “On the 20
21 hardness of being truthful.” In *2008 49th Annual IEEE Symposium on Founda-* 21
22 *tions of Computer Science*, 250–259, IEEE. [26] 22
23 23

24 Parkes, David C (2005), “Iterative combinatorial auctions.” In *Combinatorial Auc-* 24
25 *tions* (Peter Cramton, Yoav Shoham, Richard Steinberg, et al., eds.), chapter 2, 25
26 41–78, The MIT Press, Cambridge, Massachusetts. [3] 26

27 Picard, Jean-Claude (1976), “Maximal closure of a graph and applications to com- 27
28 binatorial problems.” *Management Science*, 22(11), 1268–1272. [21] 28
29 29

30 Rothkopf, Michael H, Aleksandar Pekeč, and Ronald M Harstad (1998), “Com- 30
31 putationally manageable combinational auctions.” *Management Science*, 44(8), 31
32 1131–1147. [5] 32

Roughgarden, Tim and Inbal Talgam-Cohen (2015), “Why prices need algorithms.” In *Proceedings of the sixteenth ACM conference on economics and computation*, 19–36. [5]

Shapley, Lloyd S and Martin Shubik (1971), “The assignment game i: The core.” *International Journal of Game Theory*, 1(1), 111–130. [33]

Thapper, Johan and Stanislav Živný (2016), “The complexity of finite-valued csp’s.” *Journal of the ACM (JACM)*, 63(4), 1–33. [4, 16]

Von Neumann, John (1950), “A certain zero-sum two-person game equivalent to the optimal assignment problem.” *Contributions to the Theory of Games*, (24), 5. [2]

Živný, Stanislav (2009), *The complexity and expressive power of valued constraints*. Ph.D. thesis, Oxford University, UK. [20, 33, 34]

Co-editor [Name Surname; will be inserted later] handled this manuscript.