Why Prices Need Algorithms

TIM ROUGHGARDEN, Stanford University
INBAL TALGAM-COHEN, Stanford University

Understanding when equilibria are guaranteed to exist is a central theme in economic theory, seemingly unrelated to computation. This paper shows that the existence of pricing equilibria is inextricably connected to the computational complexity of related optimization problems: demand oracles, revenue-maximization, and welfare-maximization. This relationship implies, under suitable complexity assumptions, a host of impossibility results. We also suggest a complexity-theoretic explanation for the lack of useful extensions of the Walrasian equilibrium concept: such extensions seem to require the invention of novel polynomial-time algorithms for welfare-maximization.

Categories and Subject Descriptors: F.2 [**Theory of Computation**]: Analysis of Algorithms and Problem Complexity; J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

Additional Key Words and Phrases: Market Equilibrium, Computational Game Theory, Equilibrium Computation, Market Design, Complexity

1. INTRODUCTION

Computational complexity has already had plenty to say about economic equilibria. Some of the most celebrated results in algorithmic game theory concern equilibrium concepts that have guaranteed existence, like Nash equilibria in finite games, or market equilibria in markets with divisible goods and concave utilities. These results determine whether or not such equilibria can be computed by polynomial-time algorithms under standard complexity assumptions (see, e.g., [Chen et al. 2009b; Daskalakis et al. 2009; Chen et al. 2009a]). Computational complexity is informative also for equilibrium concepts that do not have guaranteed existence, such as pure Nash equilibria in concise games, or market equilibria in markets with concave production functions (see, e.g., [Fischer et al. 2006; Papadimitriou and Wilkens 2011]). For example, by proving that it is computationally hard to compute whether or not a given game or market admits an equilibrium of a desired type, one provides evidence that there is no "nice characterization" of the instances in which such an equilibrium exists.

The primary theme of this work is that *computational complexity can also be used to study the equilibrium existence question itself*, in that non-existence results can be derived from the computational intractability of related optimization problems, under suitable complexity assumptions.

We explore this theme in the classic setting of market-clearing prices for markets of m indivisible items, where there are n consumers and each consumer i has a valuation $v_i(S)$ for each bundle S of items. There is a large literature on understanding, for var-

---

This work is supported in part by NSF Award CCF-1215965 and the Hsieh Family Stanford Interdisciplinary Graduate Fellowship. See the Acknowledgements section before REFERENCES.
Authors' addresses: T. Roughgarden and I. Talgam-Cohen, Department of Computer Science, Stanford University; email: {tim,italgam}@cs.stanford.edu.

EC'15, June 15–19, 2015, Portland, OR, USA. ACM 978-1-4503-3410-5/15/06 ...$15.00.

http://dx.doi.org/10.1145/2764468.2764515

ious classes $\mathcal{V}$ of allowable valuations, what types of prices — item prices, anonymous bundle prices, etc. — are necessary and sufficient to guarantee the existence of pricing equilibria when all consumers have valuations in $\mathcal{V}$ (see, e.g., [Parkes and Ungar 2000; Sun and Yang 2006; Candogan et al. 2015; Ben-Zwi et al. 2013; Sun and Yang 2014; Candogan et al. 2014; Candogan and Pekec 2014]). The goal of this paper is to formalize a general version of this question, show that the answer is inextricably linked to the computational complexity of well-studied problems (like demand oracles, revenue-maximization, or welfare-maximization), and derive from this connection a number of results for the (non-)existence of pricing equilibria.

1.1. Some Highlights

While the main point of this paper is its methodology, many of our specific results are also of independent interest. We now informally describe a few of them; see later sections for the relevant formal definitions.

(1) Whenever the welfare-maximization problem for a valuation class $\mathcal{V}$ is strictly computationally harder than the demand problem for $\mathcal{V}$ (with item prices), Walrasian equilibria need not exist (Proposition 2.4). For example, for budget-additive valuations, assuming $\mathrm{P} \neq \mathrm{NP}$, the welfare-maximization problem (which is strongly NP-hard) is strictly harder than the demand problem (which is weakly NP-hard), thus ruling out the possibility of guaranteed existence of Walrasian equilibria (Corollary 2.2).

(2) Walrasian equilibria, which employ only anonymous item prices, are guaranteed to exist when consumers only want one item (unit-demand). It follows from our work that they need not exist when consumers only want a pair of items, but what if we use a richer set of prices, defined on both items and item pairs? It is easy to see that *non-anonymous* prices on item pairs recover the guaranteed existence of pricing equilibria (Observation 5.2), but are anonymous prices also sufficient? Under the assumption that $\mathrm{NP} \not\subseteq \mathrm{coNP}$, our general results imply a negative answer to this question (Corollary 4.2). This conditional non-existence stems from the facts that revenue-maximization with such prices is polynomial-time solvable, the demand problem with such valuations and prices is polynomial-time solvable, and the welfare-maximization problem with these valuations is NP-hard (Proposition 4.5).

(3) Walrasian equilibria are remarkable in that, despite using only m-dimensional prices (one per item), they are guaranteed to exist for a valuation class with dimension exponential in m (gross substitutes valuations, see Lemma 5.1). Despite much research on pricing equilibria for various valuation and pricing classes, no generalizations of Walrasian equilibria with these properties (succinctness and guaranteed existence) have been found to date. Our methodology provides an explanation by identifying an algorithmic barrier to such results: it would require a novel polynomial-time algorithm for the welfare-maximization problem, beyond solving the standard configuration linear programming relaxation of the problem (Section 5.3).

1.2. Related Work

Our study of equilibrium existence for different classes of valuations is related to the large literature on how the class of valuations shapes computational aspects of combinatorial auctions and markets [Cramton et al. 2006]. We mention here three such aspects. Additional related work relevant to Sections 4 to 6 appears in these sections.

(1) Communication: Nisan and Segal [2006] study communication aspects of welfare-maximization in market settings. Generalizing a result of [Parkes 2002], they

prove for many classes of valuations (such as submodular valuations) that welfare-maximization requires exponential communication. In particular, it requires communication of a price system that is exponential in the number of items m.

(2) Approximation: Algorithmic aspects of approximate welfare-maximization have been extensively studied, especially for the "complement-free" valuation hierarchy [Blumrosen and Nisan 2007]. Recent research expands upon this by studying valuation classes with "limited" complements and good approximation guarantees [Abraham et al. 2012; Feige et al. 2014].[1]

(3) Representation and elicitation: Both succinctness ("compactness") of valuation classes and their learnability have been studied (see, e.g., [Boutilier et al. 2004; Zinkevich et al. 2003]), and for classes of non-succinct valuations, simple sketches have been pursued ([Cohavi and Dobzinski 2014] and references within).

1.3. Organization

We begin with a discussion in Section 2 of the basic market equilibrium notion of Walrasian equilibrium. We then describe our general formalism in Section 3. Anonymous pricing is discussed in Section 4, compressed pricing is discussed in Section 5, and linear pricing is discussed in Section 6. Section 7 summarizes.

2. WALRASIAN EQUILIBRIUM

This section shows that, even for the exhaustively-studied Walrasian equilibrium concept, computational complexity is a useful tool for (conditionally) ruling out existence in many scenarios. While simple to prove, these results are conceptually interesting, and also develop intuition for our subsequent results about other types of pricing equilibria. To keep this section brief we keep formal definitions to a minimum; see Section 3 for precise explanations of all terms. We begin with an overview of the results deferring formal statements to Section 2.2.

2.1. Overview

Recall that a Walrasian equilibrium is an allocation of items to consumers together with item prices such that (1) every consumer is allocated a bundle in his *demand set*; and (2) the market clears, meaning every unallocated item has price 0. For condition (1), recall that the demand set of a consumer given prices is the set of all bundles S that maximize his payoff — his value $v(S)$ for the bundle minus his total payment for it $p(S)$. The *demand problem* for a valuation class $\mathcal{V}$ is to compute, given prices, a bundle of items in the demand set of a consumer with $v \in \mathcal{V}$. For condition (2), note an equivalent condition that is useful later when we discuss more general types of prices: that the allocation maximizes the seller's revenue (given the prices).

A Walrasian equilibrium may or may not exist in a market — it depends on the structure of the consumers' valuations. At first blush, it might seem that the main results of Gul and Stacchetti [1999] and Milgrom [2000] tell us everything we want to know about the existence question: markets with gross substitutes valuations (Definition A.1) always possess Walrasian equilibria, and for every valuation class $\mathcal{V}$ that contains all unit-demand valuations (Definition A.2) and a non-gross-substitutes valuation, there exists a market with valuations in $\mathcal{V}$ and no Walrasian equilibrium. There are, however, natural valuation classes $\mathcal{V}$ that do not include all unit-demand valuations (many examples appear below). Research on such valuation classes, to which the non-existence results of Gul and Stacchetti [1999] and Milgrom [2000] do not apply, has progressed in a relatively ad hoc fashion, relying on explicit constructions to rule

[1]Note that in our work the "baseline" valuation class — beyond which valuations are considered as exhibiting complements — is that of gross substitutes.

out the existence of Walrasian equilibria in various cases (see, e.g., [Candogan et al. 2015; Ben-Zwi et al. 2013; Sun and Yang 2014; Candogan et al. 2014; Candogan and Pekec 2014]). Is a more systematic approach possible?

Let the *allocation problem* for a valuation class $\mathcal{V}$ be to maximize welfare given a market with a valuation profile from $\mathcal{V}$. Our main conceptual point in this section is:

PROPOSITION 2.1 (INFORMAL). *A necessary condition for the guaranteed existence of a Walrasian equilibrium in markets with valuations from class $\mathcal{V}$ is that the demand problem is at least as hard computationally as the allocation problem for $\mathcal{V}$.*

Proposition 2.1 is conceptually interesting because it links a basic economic question (existence of Walrasian equilibria), which is defined without reference to computation, to computational complexity considerations. We show in Section 2.2 how it can be used to re-derive many known non-existence results in an arguably more systematic way, rooted in the mature understanding in theoretical computer science of the computational complexity of various problems. We highlight in particular two applications, following directly from the complexity of well-known computational problems under the widely believed $\mathrm{P} \neq \mathrm{NP}$ assumption — one for the class of budget-additive valuations and one for the class of positive graphical valuations (for definitions of these classes see the proofs of Corollaries 2.2 and 2.3 in Section 2.2, and also, respectively, [Lehmann et al. 2006] and [Conitzer et al. 2005]).

COROLLARY 2.2. *Assuming* $\mathrm{P} \neq \mathrm{NP}$*, there exists a market with budget-additive valuations for which there is no Walrasian equilibrium.*

COROLLARY 2.3. *Assuming* $\mathrm{P} \neq \mathrm{NP}$*, there exists a market with positive graphical valuations for which there is no Walrasian equilibrium.*

2.2. Results

In this section we formalize Proposition 2.1 and prove some applications. The proposition relies on a well-known connection between the allocation problem of maximizing welfare and between Walrasian equilibria.

PROPOSITION 2.4. *Consider a valuation class $\mathcal{V}$. If for every market with valuations from $\mathcal{V}$ there exists a Walrasian equilibrium, then the allocation problem reduces in polynomial time to the demand problem.*

PROOF. (Sketch.) Recall that the allocation problem has a canonical linear programming formulation, known as the *configuration LP* (see, e.g., [Blumrosen and Nisan 2007], Section 11.3.1). A Walrasian equilibrium exists if and only if the configuration LP has an optimal integral solution ([Blumrosen and Nisan 2007], Theorem 11.13). By assumption of Walrasian equilibrium existence, if we can solve the LP in polynomial time then we can solve the allocation problem in polynomial time. Since the dual of the configuration LP has exponentially many constraints and polynomially many variables, it is solvable using the ellipsoid method by polynomially many applications of a separation oracle. The dual constraints correspond to verifying, given item prices, that the seller is maximizing his revenue and that each consumer's allocation is in his demand set. Revenue maximization with item prices is equivalent to market clearance and thus verifiable. We conclude that the separation oracle, and thus the allocation problem, reduces to solving the demand problem. This completes the proof. □

A valuation class has *demand oracle access* if it is assumed that the demand problem can be solved by a computationally efficient oracle. The following is an immediate corollary of Proposition 2.4 for such classes.

COROLLARY 2.5. *Consider a valuation class $\mathcal{V}$ with demand oracle access. If for every market with valuations from $\mathcal{V}$ there exists a Walrasian equilibrium, then the allocation problem for $\mathcal{V}$ can be solved using a polynomial amount of computation and demand queries.*

To derive the applications described above in Corollaries 2.2 and 2.3, our main interest is in the contrapositive of Proposition 2.4 — if the allocation problem cannot be reduced to the demand problem (assuming $\mathcal{P} \neq \mathrm{NP}$), then a Walrasian equilibrium does not exist for every market. We use this to sketch the proof of Corollary 2.2 (details appear in the full version), and to prove Corollary 2.3.

PROOF OF COROLLARY 2.2 FOR BUDGET-ADDITIVE VALUATIONS — SKETCH. Recall that a budget-additive valuation v assigns values $\{v_j\}_j$ to the items and has a budget b; the value $v(S)$ of a bundle S is the aggregate value of the items capped by the budget, i.e., $v(S) = \min\{\sum_{j\in S} v_j, b\}$. The first step of the proof is to show that the demand problem for a budget additive valuation v given item prices can be solved in *pseudo-polynomial* time. The second step of the proof is to show that the allocation problem is *strongly* NP-hard, meaning NP-hard even for polynomially-bounded budget-additive valuations. We show this by reduction from the strongly NP-hard *bin packing* problem. Therefore, assuming $\mathrm{P} \neq \mathrm{NP}$, the allocation problem cannot be reduced to the demand problem, and so by Proposition 2.4 a Walrasian equilibrium is not guaranteed. This completes the proof. $\square$

PROOF OF COROLLARY 2.3 FOR POSITIVE GRAPHICAL VALUATIONS. Recall that a positive graphical valuation v is represented by a graph $G = (M, E)$ with the set of items M as vertices. The vertices and edges of G are weighted by a non-negative weight function $w(\cdot)$. The value of a bundle S is $v(S) = w(G(S))$, where $G(S)$ is the subgraph induced by the vertices in S, and $w(G(S))$ is the total weight of the subgraph's vertices and edges. The fact that the demand problem given item prices is solvable in polynomial time was observed by [Abraham et al. 2012] (Proposition 5.1): The valuation v defined by the positive graph is supermodular, and hence so is the consumer's utility function after subtracting item prices from valuation v; maximizing supermodular functions can be done in polynomial time. On the other hand, the allocation problem is NP-hard by a reduction of [Conitzer et al. 2005] (Theorem 6) from the problem of *exact cover by 3-sets*. This completes the proof. $\square$

We can also apply Proposition 2.4 to the following valuation classes to show that, assuming $\mathrm{P} \neq \mathrm{NP}$, a Walrasian equilibrium need not exist: graphical valuations with an underlying tree graph and sign consistent weights [Candogan et al. 2015]; XOS valuations with a sub-polynomial number of clauses; and other hard cases of succinct supermodular valuations. See the full version for details.

3. GENERAL FORMALISM

Walrasian equilibria do not exist for many important valuation classes. A natural idea is to permit prices that are somewhat more complex than anonymous item prices, while retaining as many of the nice properties of Walrasian equilibria as possible. This section introduces the formalism needed to evaluate the prospects of this idea. We define abstract pricing functions, prove that the first and second welfare theorems continue to hold for them, and discuss our requirement of efficient verification. Sections 4 and 5 build on the concepts of this section to rule out the guaranteed existence of any nontrivial generalization of Walrasian equilibria for many valuation classes.

3.1. Valuations and Pricings

3.1.1. Definitions. A *market* is a set M of m items sold by a seller to a set N of n consumers. An *allocation* $\vec{S} = (S_1, \ldots, S_n)$ of the items is a partial partition of the item set M among the n consumers (items are allowed to remain unallocated).

Every consumer i has a *valuation* function v_i that maps bundles of items $S \subseteq M$ to their values in $\mathbb{R}$.[2] As is standard, valuation functions are assumed to be normalized ($v_i(\emptyset) = 0$) and monotone ($v_i(S) \leq v_i(T)$ for every $S \subseteq T$), and thus non-negative. A *valuation profile* v is a set of n valuations. The *welfare* of an allocation $\vec{S}$ given a valuation profile v is $\sum_i v_i(S_i)$.

We can describe general pricings in exactly the same way we describe general valuations. Formally, a *pricing profile* p is a set of n *pricing* functions, one function p_i for each consumer, each mapping bundles of items to prices in $\mathbb{R}^+$. An *anonymous* pricing profile has the same pricing for every consumer. The *revenue* of an allocation $\vec{S}$ given a pricing p is $\sum_i p_i(S_i)$.

Consider consumer i with valuation v_i and pricing p_i. The *payoff* (also known as *utility*) of this consumer from being allocated bundle S_i is quasi-linear, i.e., is equal to $v_i(S_i) - p_i(S_i)$. His *demand set* is the family of all bundles that maximize his payoff.

We are interested in *classes* $\mathcal{V}$ ($\mathcal{P}$) of valuations (pricings). We say that a profile belongs to a class if all its valuations (pricings) belong to that class.

3.1.2. Representation. A naïve representation of valuations and pricings is of exponential size and hence computationally uninteresting. One standard way to circumvent this is via *oracle access*. We say that a valuation or pricing class has oracle access of a certain kind if its functions are computed by such oracles, whose representation and queries are considered computationally efficient. The two most common kinds of oracles are as follows.

(1) A *value oracle* represents a valuation; it gets a bundle S and returns its value. Similarly, a *price oracle* represents a pricing; it gets a bundle S and returns its price. Unless otherwise noted, we consider only valuation and pricing classes that have such oracle access.
(2) A *demand oracle* represents a valuation and is defined with respect to a pricing class $\mathcal{P}$; it gets a pricing from $\mathcal{P}$ (represented by a price oracle), and returns a bundle S in the demand set given this pricing.

Of special interest are valuations (pricings) that belong to *succinct* classes, i.e., whose value (price) oracle has an explicit description polynomial in m that also runs in polynomial time. It is hard to imagine using non-succinct pricings, except in markets with a very small number of items. Similarly, it is natural to assume that actual consumers can be modeled faithfully with succinct valuations. Most of the valuation classes we deal with in this paper are succinct.

3.2. The Allocation, Demand and Revenue Problems and their Complexity

3.2.1. The Problems. The *allocation* (*welfare-maximization*) *problem* for a valuation class $\mathcal{V}$ is defined as follows: The optimization version gets as input a market with valuations from $\mathcal{V}$, and outputs a welfare-maximizing allocation. The decision version gets an additional input $w \in \mathbb{R}$, and decides whether or not there exists an allocation with welfare at least w.[3] The *symmetric* allocation problem is a special case of the allocation problem where all consumers on the market have the same valuation.

[2]Throughout, all numerical values are assumed to have a polynomial representation in the parameter m.
[3]We will sometimes be informal about distinguishing between optimization and decision versions of a computational problem.

The *demand* (*utility-maximization*) *problem* for a valuation class $\mathcal{V}$ and pricing class $\mathcal{P}$ is defined as follows: The optimization version gets as input a valuation v from $\mathcal{V}$ and a pricing p from $\mathcal{P}$, and outputs a bundle in the demand set of a consumer with valuation v given pricing p. The decision version gets an additional input $u \in \mathbb{R}$, and decides whether or not the utility of the demand set bundles is at least u.

The *revenue-maximization* problem for a pricing class $\mathcal{P}$ is equivalent to the allocation problem for the valuation class $\mathcal{V} = \mathcal{P}$. The problem is to decide, given a pricing profile from $\mathcal{P}$ and a target revenue $r \in \mathbb{R}$, whether or not there exists an allocation with revenue at least r. The *symmetric* revenue-maximization problem is the special case where all pricings are equal.

3.2.2. Complexity. This paper focuses on three of the most fundamental computational complexity classes: P, NP and coNP. We refer the non-computer-scientist reader to the full version of our paper for short informal descriptions, or to [Arora and Barak 2009] for a detailed exposition. Consider a succinct valuation class $\mathcal{V}$. The allocation problem (and hence also the revenue-maximization problem when $\mathcal{P}$ is succinct) is in NP since for every market with valuations from $\mathcal{V}$ and a target welfare w, it can be verified in polynomial time with value queries that an allocation's welfare is $\geq w$. Similarly, the demand problem is also in NP. The allocation problem is in coNP if for every such market there is a polynomial-sized *certificate* (also known as *proof*), which verifies in polynomial time with value queries that the optimal welfare is $< w$. We do not know of a natural valuation class for which the allocation problem is known to be in coNP but not known to be in P.

3.3. Pricing Equilibrium and Verifiability

3.3.1. Definitions and Basic Properties. Consider a pricing class $\mathcal{P}$. A *pricing equilibrium* for a market is an allocation $(S_1, \ldots, S_n)$ together with a *supporting* pricing profile that belongs to $\mathcal{P}$. We say a pricing profile is supporting if the following two conditions hold:

(1) The allocation maximizes the consumers' payoffs given the pricings; in other words, for every consumer i, bundle S_i is in his demand set.
(2) The allocation maximizes the seller's revenue given the pricings.

The latter condition can also be stated symmetrically to the former one, by defining the seller's "valuation" v_0 to be zero for every bundle, and then requiring that the allocation maximize his payoff.[4] We say that pricing class $\mathcal{P}$ is supporting with respect to valuation class $\mathcal{V}$ if for every market with a valuation profile from $\mathcal{V}$ there is a supporting pricing profile from $\mathcal{P}$.

A Walrasian equilibrium is precisely a pricing equilibrium supported by an anonymous profile from the class of item pricings.

The usefulness of the generalized pricing equilibrium notion stems from the following generalization of the classic welfare theorems (proofs appear in the full version).

THEOREM 3.1 (GENERALIZED WELFARE THEOREMS).

(*1*) *Every pricing equilibrium maximizes welfare.*
(*2*) *For every welfare-maximizing allocation there exists a supporting pricing profile; moreover, every supporting pricing profile supports every welfare-maximizing allocation.*

[4] Algorithmically, however, the seller's problem of maximizing revenue is different from the consumers' problem of maximizing payoff by finding a bundle in demand — the seller must find a way to allocate the items among the consumers such that their total payment according to the pricing profile is maximized.

3.3.2. Verifiability. For a pricing equilibrium concept to be meaningful, it should be efficiently recognizable (in the computational sense) by all parties:[5] given an alleged pricing equilibrium, each consumer should be able to verify efficiently that he is given a utility-maximizing bundle (given his pricing), and the seller should be able to verify efficiently that the proposed allocation maximizes his revenue (given the pricings). For example, with gross substitutes valuations and item prices, Walrasian equilibria can be verified efficiently: the revenue-maximization condition is trivial to check (all unsold items have zero price), and each utility-maximization condition can be checked with a polynomial number of value queries [Bertelsen 2004].

Assume that the pricing class $\mathcal{P}$ is succinct and revenue-maximization is tractable. There are then two levels of verifiability. We say that pricing equilibria for $\mathcal{P}$ and valuation class $\mathcal{V}$ are *verifiable with demand oracle access* if the valuations in $\mathcal{V}$ have demand oracle access with respect to the pricings in $\mathcal{P}$. In this case, all the pricing equilibrium conditions can be verified using a polynomial amount of computation and demand queries. If, in addition, a pricing equilibrium for $\mathcal{P}$ and $\mathcal{V}$ is guaranteed to exist, then we say that the allocation problem is verifiable with demand oracle access, since the equilibrium serves as a succinct and verifiable certificate.

A stronger requirement on $\mathcal{V}$ and $\mathcal{P}$ is that pricing equilibria are *verifiable with value oracle access*. This is the case if the demand problem can be solved using a polynomial amount of computation and value queries. When $\mathcal{P}$ and $\mathcal{V}$ guarantee the existence of a pricing equilibrium we again say that the allocation problem is verifiable with value oracle access. Finally, if $\mathcal{V}$ is also succinct, the pricing equilibria are verifiable in polynomial time and the allocation problem belongs to coNP.

4. ANONYMOUS PRICING

Walrasian equilibria employ item prices that are "simple" in three respects: they are anonymous (common to all consumers); they are succinct; and they make the revenue-maximization problem tractable. This section and the next study, in general, when we can and cannot obtain the first and second property, respectively, while keeping the third property as a constraint.

4.1. Related Work

Consider a valuation class and its supporting pricing classes. It is helpful to classify the pricing classes into two categories: In the first category, pricings belong to the same class as the valuations, and in the second category, pricings are untethered from the valuation class and are allowed to belong to a broader or altogether different class. While the second category may be required in order to achieve anonymity, keeping supporting prices simple is important for verifiability of the resulting pricing equilibria.

There are two main examples in the literature of valuation classes supported by anonymous pricing profiles.

(1) *Gross substitutes and anonymous item prices*: This canonical example belongs to the first category, since item prices correspond to additive valuations, which belong to the class of gross substitutes.
(2) *Supperadditive valuations and anonymous bundle prices*: A valuation v is superadditive if $v(S \cup S') \geq v(S) + v(S')$ for every disjoint bundles $S, S' \subseteq M$ [Parkes and Ungar 2000; Vohra 2011; Sun and Yang 2014]. This example belongs to the

[5]One could argue that the stronger requirement of polynomial-time computability should also hold. Since our main results are negative, adopting the weaker requirement of efficient verifiability only strengthens our results.

second category, since bundle prices correspond to general valuations, which are a strict superclass of superadditive valuations. Even more is known about anonymous bundle prices: Bikhchandani and Ostroy [2002] formulate a linear program that gives an instance by instance characterization of markets for which such supporting prices exist, and Parkes and Ungar [2000] develop an algorithm for finding them. They further show that anonymous bundle prices support certain XOR valuation subclasses.

Candogan [2013] studies anonymous graphical pricings, and establishes the following negative finding: a valuation profile from the class of graphical valuations based on series-parallel graphs is supported by an anonymous graphical pricing profile if and only if it is also supported by anonymous item prices (i.e., a Walrasian equilibrium exists for this market instance).

4.2. Overview

In this overview section we present our main findings on anonymous pricing profiles and derive applications; details appear in Section 4.3. We focus on succinct pricing classes for which the symmetric revenue-maximization problem is tractable.[6] We establish a similar condition to that of Proposition 2.1, where we showed that a necessary condition for the existence of a Walrasian equilibrium is that the demand problem is as hard as the allocation one.

PROPOSITION 4.1 (INFORMAL). *Consider a succinct pricing class $\mathcal{P}$ for which the symmetric revenue-maximization problem is tractable. A necessary condition for the guaranteed existence of a pricing equilibrium with an anonymous pricing profile from $\mathcal{P}$ in markets with valuations from class $\mathcal{V}$ is that the demand problem is at least as hard computationally as the allocation problem for these classes.*

We demonstrate three applications of Proposition 4.1.

We begin by defining *hypergraph pricings* — a succinct pricing class for which the symmetric revenue-maximization problem is tractable: Recall the class of positive graphical valuations from Corollary 2.3. A natural generalization is to consider succinct hypergraphs in place of graphs.[7] By definition, the resulting pricing class is succinct. The symmetric revenue-maximization problem is tractable since revenue is maximized by allocating all items to a single consumer.

As a first application, consider a natural generalization of unit-demand valuations (Definition A.2) — *pair-demand* valuations, where there exist values $\{v_{i,j}\}_{i,j}$ for item pairs such that for every bundle S, $v(S) = \max_{i,j\in S}\{v_{i,j}\}$. Since anonymous item prices support unit-demand valuations, one could hope that anonymous hypergraphical prices support pair-demand valuations. Proposition 4.1 refutes this hope.

COROLLARY 4.2. *Assuming* $\mathrm{NP} \not\subseteq \mathrm{coNP}$*, there exists a market with pair-demand valuations for which there is no pricing equilibrium supported by an anonymous profile of hypergraph pricings.*

As a second application, we return to positive graphical valuations, which form a succinct subclass of supermodular valuations and as such may be considered a natural candidate for the existence of a verifiable anonymous pricing equilibrium. However:

[6]The previous "success stories" of anonymous pricing profiles beyond item prices – anonymous bundle prices supporting supermodular and other valuation classes – escape our impossibility results in this section by being non-succinct.

[7]Such a pricing is represented by a hypergraph over the set of items M as vertices, with non-negative weights on the vertices and hyperedges. The price of a bundle S is the total weight of the hypergraph induced by it.

COROLLARY 4.3. *Assuming* $\text{NP} \not\subseteq \text{coNP}$*, there exists a market with positive graphical valuations for which there is no pricing equilibrium supported by an anonymous profile of hypergraph pricings.*

As a third application, consider a generalization of the Sun and Yang [2006] market: Sun and Yang define a market with two sets of items referred to as "tables" and "chairs", where consumers have unit-demand valuations for both tables and chairs, but consider a pair of (table, chair) to be complementary. They show that a Walrasian equilibrium is guaranteed to exist in their market. An obvious generalization is to allow *three* sets of items with cross-set complementarities (*cf.*, [Teytelboym 2014]), however:

COROLLARY 4.4. *Assuming* $\text{NP} \not\subseteq \text{coNP}$*, there exists a generalized Sun and Yang market for which there is no pricing equilibrium supported by an anonymous profile of hypergraph pricings.*

4.3. Results

We first prove a more formal version of Proposition 4.1 and then establish the applications described above.

PROPOSITION 4.5. *Consider a succinct valuation class* $\mathcal{V}$ *and a succinct pricing class* $\mathcal{P}$ *for which the symmetric revenue-maximization problem is tractable. If for every market with valuations from* $\mathcal{V}$ *there exists a pricing equilibrium with an anonymous pricing profile from* $\mathcal{P}$*, then the allocation problem belongs to* coNP *whenever the demand problem belongs to* coNP.

PROOF. Assume the demand problem belongs to coNP. Consider a market with valuations from $\mathcal{V}$. We want to show the existence of a polynomial-sized certificate, which can certify in polynomial time that the optimal welfare is $< w$. By the first welfare theorem (Theorem 3.1), every pricing equilibrium maximizes welfare, and by the proposition's assumptions, a polynomial-sized equilibrium with an anonymous pricing profile is guaranteed to exist for $\mathcal{V}$ and $\mathcal{P}$. Consider such a pricing equilibrium with welfare $< w$; we know that symmetric revenue-maximization is tractable and so verifying it reduces to verifying n instances of the demand problem, one per consumer. Since the demand problem is assumed to be in coNP, there are polynomial-sized certificates for verifying these in polynomial time. Together with the pricing equilibrium itself they form the required certificate, completing the proof. □

For valuation classes with demand oracle access, a similar proof shows the following.

COROLLARY 4.6. *Consider a valuation class* $\mathcal{V}$ *with demand oracle access and a succinct pricing class* $\mathcal{P}$ *for which the symmetric revenue-maximization problem is tractable. If for every market with valuations from* $\mathcal{V}$ *there exists a pricing equilibrium with an anonymous pricing profile from* $\mathcal{P}$*, then the allocation problem for* $\mathcal{V}$ *is verifiable with demand oracle access.*

To derive the applications described above in Corollaries 4.2 to 4.4, our main interest is in the following immediate corollary of Proposition 4.5.

COROLLARY 4.7. *Assuming* $\text{NP} \not\subseteq \text{coNP}$*, for classes* $\mathcal{V}$ *and* $\mathcal{P}$ *as in Proposition 4.5 for which the allocation problem is* NP*-hard and the demand problem is in* P*, there exists a market with valuations in* $\mathcal{V}$ *for which there is no pricing equilibrium supported by an anonymous pricing profile from* $\mathcal{P}$*.*

4.3.1. Applications

PROOF OF COROLLARY 4.2 FOR PAIR-DEMAND VALUATIONS. In the full version we show that the allocation problem for pair-demand valuations is NP-hard, by a reduction from the problem of *three-dimensional matching*. On the other hand, the demand problem for pair-demand valuations with hypergraphical prices is in P, simply by computing and comparing the payoffs from all possible pairs of items. Applying Corollary 4.7 completes the proof. □

To prove Corollaries 4.3 and 4.4, the following definition is useful. A consumer has a *triplet valuation* if there is a triplet of items j, k, l such that his value for a bundle S is as follows: $v_{j,k}$ if $j, k \in S$ and $l \notin S$; $v_{k,l}$ if $k, l \in S$ and $j \notin S$; $v_{j,l}$ if $j, l \in S$ and $k \notin S$; $v_{j,k} + v_{k,l} + v_{j,l}$ if $j, k, l \in S$; and zero otherwise. Since triplet valuations are strict subclasses of positive graphical valuations and of the valuations in the generalized Sun and Yang market, Corollaries 4.3 and 4.4 follow immediately from the next lemma.

LEMMA 4.8. *Assuming* $\mathrm{NP} \nsubseteq \mathrm{coNP}$*, there exists a market with triplet valuations for which there is no pricing equilibrium supported by an anonymous profile of hypergraph pricings.*

PROOF. The allocation problem for triplet valuations is NP-hard by a reduction from the problem of exact cover by 3-sets [Conitzer et al. 2005]. On the other hand, the demand problem for triplet valuations with hypergraphical prices is in $\mathcal{P}$, simply by computing and comparing the payoffs from every pair of items in the triplet and from the entire triplet. Applying Corollary 4.7 completes the proof. □

5. COMPRESSED PRICING

This section focuses on the low-dimensionality or *compressed* aspect of the prices used in Walrasian equilibria — there is only one price for each of the m items, and yet existence is guaranteed even for the high-dimensional class of gross substitutes valuations. When else can we achieve guaranteed existence of succinct equilibria, even for non-succinct valuations?

5.1. Related Work and Background

Compression of valuations is an important theme in mechanism and market design, with a classic trade-off between expressiveness of the valuations and simplicity of the market mechanism.

In the context of market equilibria, the class of gross substitutes stands out as the canonical example for which simple m-dimensional equilibria exist, despite the fact that the dimension of this valuation class is exponential in m.[8] Note that the high-dimensionality rules out the possibility of a compact encoding for gross substitutes (*cf.*, [Hatfield and Milgrom 2005; Hatfield et al. 2012]):

LEMMA 5.1. *The class of gross substitutes valuations is not succinct.*

Another research direction is compressed bid spaces for simple auction formats. In this context, the level of compression affects how close equilibria of these auctions can get to a welfare-maximizing allocation. See [Christodoulou et al. 2008] for item bidding in combinatorial auctions, and [Duetting et al. 2013; Babaioff et al. 2014] for recent extensions. [Feldman et al. 2015] further extend this line of work by considering posted price mechanisms.

[8]The gross substitutes class contains all weighted matroid rank functions. Even with constant-size weights and restricted to partition matroids, there are doubly-exponentially many of these.

5.2. Overview

We begin with a trivial "line in the sand:" when the pricing class is equal to the valuation class — i.e., there is no compression — a (non-anonymous) pricing equilibrium always exists.

OBSERVATION 5.2. *Consider a valuation class $\mathcal{V}$ and an identical pricing class $\mathcal{P} = \mathcal{V}$. For every market with valuations from $\mathcal{V}$ there exists a pricing equilibrium with pricings from $\mathcal{P}$.*

PROOF. Given a market with valuations from $\mathcal{V}$, the following is a pricing equilibrium: Pick a welfare-maximizing allocation, and for every consumer i let his pricing p_i be equal to his valuation v_i. The consumers are indifferent among different bundles and so trivially maximize their payoff. The revenue is equal to the welfare and so is maximized by the allocation. □

In fact, every welfare-maximizing allocation can be supported by pricings from the class $\mathcal{P} = \mathcal{V}$. If $\mathcal{V}$ is succinct, this characterizes when finding a verifiable pricing equilibrium is tractable.

COROLLARY 5.3. *Consider a succinct valuation class $\mathcal{V}$. There exists a polynomial time algorithm that finds a verifiable pricing equilibrium for every market with valuations from $\mathcal{V}$ if and only if the allocation problem for $\mathcal{V}$ is in* P.

One example of a tractable allocation problem is for the class of feature-based valuations with a constant number of features (see [Candogan and Pekec 2014] and Section 6.2).[9]

Given the above observation, the main question in this section is: how much can pricings be compressed in comparison to valuations, without losing the ability to support efficient allocations? I.e., how much can the pricing class be shrunk with respect to the valuation class while still guaranteeing the existence of a pricing equilibrium?

We first present a simple necessary condition on the valuation and pricing classes (Proposition 5.4). Like Propositions 2.4 and 4.5, it is based on computational complexity considerations. We then address, in Section 5.3, the question of whether there are classes for which partial compression is possible. We show an example with a supporting pricing class of dimension strictly smaller than the valuation class, yet strictly larger than m item prices. Section 6 revisits these questions from an instance-by-instance perspective, and uses linear programming to characterize when succinct "linear" supporting pricings exist.

PROPOSITION 5.4. *Consider a succinct valuation class $\mathcal{V}$ and a succinct pricing class $\mathcal{P}$ for which the revenue-maximization problem is tractable. If for every market with valuations from $\mathcal{V}$ there exists a pricing equilibrium with pricings from $\mathcal{P}$, then the allocation problem belongs to* coNP *whenever the demand problem belongs to* coNP.

The proof is similar to that of Proposition 4.5, and a corollary similar to Corollary 4.6 follows for valuation classes with demand oracle access.

COROLLARY 5.5. *Consider a valuation class $\mathcal{V}$ with demand oracle access and a succinct pricing class $\mathcal{P}$ for which the revenue-maximization problem is tractable. If for every market with valuations from $\mathcal{V}$ there exists a pricing equilibrium with pricings from $\mathcal{P}$, then the allocation problem for $\mathcal{V}$ is verifiable with demand oracle access.*

[9]Feature-based pricings are of interest in practical settings, where the seller often needs to set prices based on a small set of summarizing features (*cf.*, [Dughmi et al. 2014]). Tractability follows from dynamic programming, utilizing the constant number of different item *types*.

5.3. Partial Compression

5.3.1. A Rare Species. Despite the large literature on pricing equilibria for various valuation classes and pricing classes, we are unaware of any previously studied examples of classes $\mathcal{V}$ and $\mathcal{P}$ that meet the following criteria: (i) anonymous item prices are insufficient to support all markets with valuations in $\mathcal{V}$; (ii) the prices in $\mathcal{P}$ are sufficient to support all markets with valuations in $\mathcal{V}$; (iii) $\mathcal{P}$ is succinct and strictly smaller than $\mathcal{V}$ (where $\mathcal{V}$ may or may not be succinct); (iv) pricing equilibria for $\mathcal{V}$ and $\mathcal{P}$ are verifiable with demand oracle access (i.e., the revenue-maximization condition can be efficiently verified). In this sense, there are no known non-trivial generalizations of Walrasian equilibria!

Our complexity-theoretic methodology provides an explanation for the paucity of examples. Suppose the classes $\mathcal{V}$ and $\mathcal{P}$ satisfy (i)—(iv). By (ii)—(iv) and Corollary 5.5, the welfare-maximization problem for $\mathcal{V}$ is verifiable with demand oracle access, and this can be thought of as membership in coNP with demand oracle access. Because the problem can also be thought of as belonging to NP with demand oracle access, this suggests (but does not prove, of course) that the welfare-maximization problem can in fact be solved in polynomial time with demand oracle access. On the other hand, by (i), the configuration linear program fails to solve the welfare-maximization problem in polynomial time with demand oracle access (recall the proof of Proposition 2.4 and [Blumrosen and Nisan 2007], Theorem 11.13). Thus, a non-trivial generalization of Walrasian equilibria in the above sense requires a novel polynomial-time algorithm for the welfare-maximization problem (unless such a problem belongs to $(\mathrm{coNP} \cap \mathrm{NP}) \setminus \mathcal{P}$ with demand oracles)![10]

We remark that the same argument can be repeated for a succinct valuation class $\mathcal{V}$ for which pricing equilibria with $\mathcal{P}$ are verifiable.

5.3.2. An example. We show there are (somewhat contrived) classes $\mathcal{V}$ and $\mathcal{P}$ that satisfy properties (i)—(iv) above. It is no coincidence that also (v) the welfare-maximization problem for $\mathcal{V}$ can be solved in polynomial time, but not directly by the configuration LP.

Example 5.6. We define a family of markets with two consumers, whose valuations belong to a slight variation of the gross substitutes class (which remains not succinct). These markets are supported by a succinct pricing class, but not by item prices.

Define the valuation class $\mathcal{V}$ as follows: There are two special items, without loss of generality items 1 and 2. Every valuation $v \in V$ is equal to a gross substitutes valuation with the possible deviation that items 1 and 2 are allowed to complement each other, i.e., to be valued positively as a couple but as zero on their own. The following transformation of v must recover a gross substitutes valuation v': unify items 1 and 2 into a single item. For every bundle $S \subseteq M$, let S' be the same bundle after unifying items 1 and 2 (if $S = M$ let M' be the reduced item set). Observe that $v'(S') = v(S)$.

CLAIM 5.7. *There exists a succinct pricing class $\mathcal{P}$ such that for every market with two consumers whose valuations belong to the non-succinct class $\mathcal{V}$ defined above, there are supporting pricings in $\mathcal{P}$. Moreover, there exists such a market for which there is no Walrasian equilibrium.*

The proof of Claim 5.7 appears in the full version.

[10] While there are of course other algorithms (like greedy algorithms) that compute a welfare-maximizing allocation in various special cases, in all known such cases the configuration LP also solves the problem exactly.

6. LINEAR PRICING

Most of our results thus far have been impossibility results, ruling out the guaranteed existence of pricing equilibria for different classes of valuations and pricings. This section characterizes equilibrium existence on an instance-by-instance (rather than class-by-class) basis. We formulate the notion of a succinct linear valuation or pricing, and show that this notion captures and generalizes many previously studied valuation classes. We then give a linear programming characterization of the market instances supported by succinct linear pricings.

6.1. Definitions

Consider a set function f on the ground set of items M. A naïve representation is linear in the number of sets and exponential in the number of items m. We propose an alternative representation based on a set $\mathcal{L}$ of *pseudo-items*.

Definition 6.1. A *linear representation* is a mapping $L : 2^M \to 2^{\mathcal{L}}$ from bundles of items in M to bundles of pseudo-items in $\mathcal{L}$.

A linear representation can be seen as a bipartite graph with bundles of items on one side and individual pseudo-items on the other. Each bundle of items is connected to several pseudo-items. Now we can add weights to the pseudo-items. In this way a linear representation associates a value with each bundle of items: the total weight of the pseudo-items the bundle is connected to.

Definition 6.2. We say that a set function f on a ground set M has linear representation L if there exist weights $w_\ell \in \mathbb{R}$ for the pseudo-items such that for every bundle of items $S \subseteq M$, $f(S) = \sum_{\ell \in L(S)} w_\ell$.

Observe that every set function f has an exponentially large linear representation as follows: Let $\mathcal{L} = 2^M$, i.e., let there be a pseudo-item ℓ_S for every bundle of items S. For every S let $L(S) = \{\ell_S\}$ and let the corresponding weight be $w_{\ell_S} = f(S)$. Then $\sum_{\ell \in L(S)} w_\ell = f(S)$ as required. However, set functions can have multiple linear representations, and we are most interested in those that are *succinct*.

Definition 6.3. A family of linear representations is *succinct* if for every linear representation in the family, the set $\mathcal{L}$ of pseudo-items is of polynomial size in $m = |M|$, and the mapping L can be implemented by a polynomial time algorithm.

Succinct linear functions are functions that have succinct linear representations with polynomially-represented weights $\{w_\ell\}$.

6.2. Succinct Linear Valuations in the Literature

Succinct linear valuations capture many previously studied succinct valuation classes. This section includes several examples (for additional examples see the full version).[11]

General hypergraphical valuations. Consider a valuation v defined by a hypergraph (M, E) over the items, with a polynomial number of hyperedges $e \subseteq M$ with weights c_e. The following is a succinct linear representation of v: Let $\mathcal{L} = E$ and $L(S) = \{e \mid e \subseteq S\}$. Let the weights be $w_e = c_e$. Then we get $\sum_{\ell \in L(S)} w_\ell = \sum_{e \subseteq S} c_e = v(S)$, as required, and computing $L(S)$ requires polynomial time. A further generalization is *additively decomposable valuations* [Candogan 2013] (Section 5.5), and these also form a strict subset of succinct linear valuations.

[11] An example of a valuation class that is not captured is valuations that map bundles to the product of their items' values.

Explicit coverage valuations. Consider a valuation v defined by a ground set $\mathcal{G}$ and a subset $G_j \subseteq \mathcal{G}$ associated with each item j. The value $v(S)$ of every bundle S is $|\bigcup_{j \in S} G_j|$. The following is a linear representation of v: Let $\mathcal{L} = \mathcal{G}$ and $L(S) = \bigcup_{j \in S} G_j$. If all weights are set to 1 then we get $\sum_{\ell \in L(S)} w_\ell = |L(S)| = v(S)$, as required. Computing $L(S)$ if the coverage valuation is given explicitly by a bipartite graph encoding is in polynomial time in the size of the encoding.

Feature-based valuations. Consider a valuation v described by a multi-unit valuation v_f for each of polynomially many features f, and polynomial-time functions $c_f(S)$ that return the number of times feature f appears in bundle S [Candogan and Pekec 2014]. The value $v(S)$ for every bundle S is equal to $\sum_f v_f(c_f(S))$. The following is a succinct linear representation of v: Let $\mathcal{L} = \{\ell_{f,k}\}$ and $L(S) = \bigcup_f \ell_{f,c_f(S)}$. Let the weights be $w_{\ell_{f,c_f(S)}} = v_f(c_f(S))$. Then we get $\sum_{\ell \in L(S)} w_\ell = \sum_f v_f(c_f(S)) = v(S)$, as required, and computing $L(S)$ requires polynomial time.

Additional examples. (See the full version for details.)

(1) Succinct endowed assignment valuations, proposed as a succinct subclass of gross substitutes by Hatfield and Milgrom [2005].
(2) Budget-additive valuations (e.g., [Lehmann et al. 2006]).
(3) XOS valuations with polynomially many clauses, and their generalization to maximum over polynomially many positive hypergraph valuations with rank k [Feige et al. 2014].
(4) $\mathrm{GGS}(k, M)$ valuations with a constant k [Ben-Zwi et al. 2013].
(5) Sketches of valuations [Badanidiyuru et al. 2012; Cohavi and Dobzinski 2014].

6.3. Linear Program Characterization

The linear programming formulation in this section and the welfare theorems in the next apply to general (not necessarily succinct) linear representations. Given a profile of linear representations $\{L_i\}$, one for each consumer, we formulate a linear program that characterizes the existence of a pricing equilibrium with price functions represented by $\{L_i\}$. Without loss of generality we assume a unified (across consumers) set of pseudo-items $\mathcal{L}$, and define the following program $\mathrm{LP1}$ for maximizing welfare.

$$
\begin{array}{lcll}
\max & \sum_{i=1}^{n} \sum_S v_i(S) x_{iS} & & \\
s.t. & \sum_S x_{iS} \le 1 & \forall 1 \le i \le n & (1) \\
 & \sum_{i=1}^{n} \sum_{S:\ell \in L_i(S)} x_{iS} = \sum_{i=1}^{n} \sum_{\mu:\ell \in L_i(\mu_i)} \sigma_\mu & \forall \ell \in \mathcal{L} & (2) \\
 & \sum_\mu \sigma_\mu \le 1 & & (3) \\
 & x_{iS} \ge 0 & \forall 1 \le i \le n, \forall S & \\
 & \sigma_\mu \ge 0 & \forall \text{allocation } \mu &
\end{array}
$$

The variables of $\mathrm{LP1}$ are: x_{iS}, indicating whether bundle S is allocated to consumer i; and σ_μ, indicating whether ordered partition $\mu = (\mu_1, \ldots, \mu_n)$ is the chosen allocation (where consumer i gets part μ_i and all items are allocated). The number of variables of $\mathrm{LP1}$ is exponential. Constraint (1) ensures no consumer is allocated more than one bundle; Constraint (2) matches the number of times pseudo-item ℓ appears in the bundles allocated to the consumers with the number of its appearances in the allocation; Constraint (3) ensures a single allocation is chosen.

As for the dual program, the variables are: π_i, the payoff of player i (including the seller as player 0); and w_ℓ, the weight of pseudo-item ℓ. Constraints (4) and (5) ensure that π_i is indeed player i's payoff. The number of constraints is exponential.

$$
\begin{aligned}
\min \quad & \textstyle\sum_{i=0}^{n} \pi_i \\
s.t. \quad \pi_i &\geq v_i(S) - \textstyle\sum_{\ell \in L_i(S)} w_\ell \quad \forall 1 \leq i \leq n, \forall S \qquad (4)\\
\pi_0 &\geq \textstyle\sum_{i=1}^{n} \sum_{\ell \in L_i(\mu_i)} w_\ell \quad \forall \text{allocation } \mu \qquad (5)\\
\pi_i &\geq 0 \quad \forall 0 \leq i \leq n
\end{aligned}
$$

6.4. Welfare Theorems

We now state versions of the welfare theorems for LP1 and DP1 (*cf.*, [Blumrosen and Nisan 2007], Theorems 11.13 and 11.15). The proofs appear in the full version of the paper.

THEOREM 6.4 (FIRST WELFARE THEOREM). *Every pricing equilibrium whose pricing functions have linear representations $\{L_i\}$ maximizes welfare over all fractional allocations that are feasible solutions to* LP1.

THEOREM 6.5 (SECOND WELFARE THEOREM). *If an integral optimal solution that corresponds to a feasible allocation exists for* LP1*, then a pricing equilibrium whose pricing functions have linear representations $\{L_i\}$ also exists.*

PROOF. By complementary slackness — see the full version for details. □

Combining the two theorems gives the following characterization.

COROLLARY 6.6 (CHARACTERIZATION). *A pricing equilibrium whose pricings have linear representations $\{L_i\}$ exists if and only if* LP1 *has an integral optimal solution that corresponds to a feasible allocation.*[12]

6.5. Implications for Guaranteed Existence

Returning to our theme of necessary conditions for the guaranteed existence of pricing equilibria, we can use Corollary 6.6 to extend Proposition 2.4 to arbitrary succinct linear pricings.

COROLLARY 6.7. *Consider a valuation class $\mathcal{V}$ and a succinct linear pricing class $\mathcal{P}$ for which the revenue-maximization problem is tractable. If for every market with valuations from $\mathcal{V}$ there exists a pricing equilibrium with pricings from $\mathcal{P}$, then the allocation problem reduces in polynomial time to the demand problem.*

The proof follows that of Proposition 2.4, using the characterization in Corollary 6.6 in place of the classical linear programming characterization of the existence of Walrasian equilibria.

Corollary 6.7 differs from Proposition 5.4 in two respects: in the hypothesis, the pricing class $\mathcal{P}$ is assumed to be linear in addition to being succinct; and in the conclusion, the computation (rather than just the verification) of a welfare-maximizing allocation reduces to that of the demand problem. Equivalently, the two "coNP" terms of Proposition 5.4 are replaced in Corollary 6.7 by "P."

[12]The condition that the solution correspond to a feasible allocation can alternatively be encoded into the linear program or into the pricing: Adding the constraint $\sum_{i=1}^{n} \sum_{S:j \in S} x_{iS} = \sum_{i=1}^{n} \sum_{\mu : j \in \mu_i} \sigma_\mu$ for every item j to LP1 would achieve this, with additional dual item price variables added to DP1. The pricing would then be the sum of the item prices and linear prices. Another possibility is to include in $\mathcal{L}$ a pseudo-item for every item, and for every bundle S to add to the mapping $L(S)$ the pseudo-item corresponding to every $j \in S$. In this case the pricing includes item prices by definition.

7. CONCLUSION AND OPEN QUESTIONS

The well-studied problem of proving or disproving the guaranteed existence of pricing equilibria seems to have nothing to do with computation. As this paper demonstrates, however, computational complexity offers numerous insights into the problem, and provides general techniques for proving impossibility results. For example, many (conditional) non-existence results for different types of pricing equilibria (Walrasian, anonymous, compressed, etc.) follow easily from the known computational complexity of various optimization problems, obviating the need for ad hoc explicit constructions without equilibria. Similarly, this methodology demystifies the dearth of useful extensions of the Walrasian equilibrium concept, by linking the existence of such extensions to algorithmic progress on the welfare-maximization problem.

ACKNOWLEDGMENTS

The authors would like to thank Dr. Vasilis Gkatzelis of Stanford University for multiple discussions from which we greatly benefited.

REFERENCES

Ittai Abraham, Moshe Babaioff, Shaddin Dughmi, and Tim Roughgarden. 2012. Combinatorial Auctions with Restricted Complements. In *Proceedings of the 13th ACM Conference on Economics and Computation*.

Sanjeev Arora and Boaz Barak. 2009. *Computational Complexity: A Modern Approach*. Cambridge University Press.

Moshe Babaioff, Brendan Lucier, Noam Nisan, and Renato Paes Leme. 2014. On the efficiency of the Walrasian mechanism. In *Proceedings of the 15th ACM Conference on Economics and Computation*.

Ashwinkumar Badanidiyuru, Shahar Dobzinski, Hu Fu, Robert Kleinberg, Noam Nisan, and Tim Roughgarden. 2012. Sketching valuation functions. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms*. 1025–1035.

Oren Ben-Zwi, Ron Lavi, and Ilan Newman. 2013. Ascending Auctions and Walrasian Equilibrium. (2013). Working paper.

Alejandro Bertelsen. 2004. *Substitutes Valuations and M-sharp Concavity*. Master's thesis. Hebrew University of Jerusalem.

Sushil Bikhchandani and Joseph M. Ostroy. 2002. The package assignment model. *Journal of Economic Theory* 107, 2 (2002), 377–406. Following 1998 technical report.

Liad Blumrosen and Noam Nisan. 2007. Combinatorial Auctions. In *Algorithmic Game Theory*, Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani (Eds.). Chapter 11.

Craig Boutilier, Ronen I. Brafman, Carmel Domshlak, Holger H. Hoos, and David Poole. 2004. CP-nets: A Tool for Representing and Reasoning with Conditional Ceteris Paribus Preference Statements. *Journal of Artificial Intelligence Research* 21 (2004), 135–191.

Ozan Candogan. 2013. *Dynamic Strategic Interactions: Analysis and Mechanism Design*. Ph.D. Dissertation. Massachusetts Institute of Technology.

Ozan Candogan, Asuman Ozdaglar, and Pablo Parrilo. 2014. Iterative Auction Design for Graphical Valuations Part II: General Graphs. (2014). In submission.

Ozan Candogan, Asuman Ozdaglar, and Pablo Parrilo. 2015. Iterative Auction Design for Tree Valuations. (2015). To appear in Operations Research.

Ozan Candogan and Sasa Pekec. 2014. Efficient Iterative Auctions for Multi-Featured Items: A Network Flow Approach. (2014). In submission.

Xi Chen, Decheng Dai, Ye Du, and Shang-Hua Teng. 2009a. Settling the Complexity of Arrow-Debreu Equilibria in Markets with Additively Separable Utilities. In *Proceedings of the 50th Symposium on Foundations of Computer Science*. 273–282.

Xi Chen, Xiaotie Deng, and Shang-Hua Teng. 2009b. Settling the complexity of computing two-player Nash equilibria. *J. ACM* 56, 3 (2009).

George Christodoulou, Annamria Kovacs, and Michael Schapira. 2008. Bayesian Combinatorial Auctions. In *Proceedings of the 35th International Colloquium on Automata, Languages and Programming*.

Keren Cohavi and Shahar Dobzinski. 2014. Faster and Simpler Sketches of Valuation Functions. (2014). Working paper.

Vincent Conitzer, Tuomas Sandholm, and Paolo Santi. 2005. Combinatorial Auctions with k-wise Dependent Valuations. In *Proceedings of the 20th AAAI Conference on Artificial Intelligence*.

Peter C. Cramton, Yoav Shoham, and Richard Steinberg. 2006. *Combinatorial auctions*.

Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. 2009. The Complexity of Computing a Nash Equilibrium. *SIAM J. Comput.* 39, 1 (2009), 195–259.

Paul Duetting, Monika Henzinger, and Martin Starnberger. 2013. Valuation Compressions in VCG-Based Combinatorial Auctions. In *Proceedings of the 9th International Workshop on Internet and Network Economics*.

Shaddin Dughmi, Nicole Immorlica, and Aaron Roth. 2014. Constrained Signaling in Auction Design. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms*.

Uriel Feige, Michal Feldman, Nicole Immorlica, Rani Izsak, Brendan Lucier, and Vasilis Syrgkanis. 2014. A Unifying Hierarchy of Valuations with Complements and Substitutes. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*.

Michal Feldman, Nick Gravin, and Brendan Lucier. 2015. Combinatorial Auctions via Posted Prices. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms*.

Felix A. Fischer, Markus Holzer, and Stefan Katzenbeisser. 2006. The influence of neighbourhood and choice on the complexity of finding pure Nash equilibria. *Information Prrocessing Letters* 99, 6 (2006), 239–245.

Faruk Gul and Ennio Stacchetti. 1999. Walrasian Equilibrium with Gross Substitutes. *Journal of Economic Theory* 87 (1999), 95–124.

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

John W. Hatfield and Paul R. Milgrom. 2005. Matching with contracts. *American Economic Review* 95 (2005), 913–935.

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

Paul R. Milgrom. 2000. Putting Auction Theory to Work: The Simultaneous Ascending Auction. *Journal of Political Economy* 108, 2 (2000), 245–272.

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

Christos H. Papadimitriou and Christopher A. Wilkens. 2011. Economies with non-convex production and complexity equilibria. In *Proceedings of the 12th ACM Conference on Economics and Computation*. 137–146.

David C. Parkes. 2002. Price-based information certificates for minimal-revelation combinatorial auctions. In *Proceedings of the 1st International Joint Conference on Autonomous Agents and Multi-Agent Systems*. 103–122.

David C. Parkes and Lyle H. Ungar. 2000. Iterative Combinatorial Auctions: Theory and Practice. In *Proceedings of the 17th AAAI Conference on Artificial Intelligence*. 74–81.

Ning Sun and Zaifu Yang. 2006. Equilibria and indivisibilities: Gross substitutes and complements. *Econometrica* 74, 5 (2006), 1385–1402.

Ning Sun and Zaifu Yang. 2014. An efficient and incentive compatible dynamic auction for multiple complements. *Journal of Political Economy* 122, 2 (2014), 422–488.

Alex Teytelboym. 2014. Gross substitutes and complements: a simple generalization. *Economics Letters* 123, 2 (2014), 135–138.

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

Martin A. Zinkevich, Avrim Blum, and Tuomas Sandholm. 2003. On Polynomial Time Preference Elicitation with Value Queries. In *Proceedings of the 4th ACM Conference on Economics and Computation*.

A. APPENDIX

Definition A.1. A valuation v is *gross substitutes* if for every two vectors of item prices $\vec{p}, \vec{q}$ such that $\vec{q} \geq \vec{p}$, for every bundle S in the demand set of v given $\vec{p}$, there exists a bundle T in the demand set of v given $\vec{q}$ which contains every item $j \in S$ whose price according to $\vec{q}$ equals its price according to $\vec{p}$.

Definition A.2. A valuation v is *unit demand* if there are item values $\{v_j\}_j$ such that for every bundle S, $v(S) = \max_{j \in S}\{v_j\}$.