TWO MEASURES OF THE DEGREE OF CONSUMER RATIONALITY VIA GRAPH PROPERTIES

Shant Boodaghians

Department of Mathematics and Statistics

McGill University, Montreal

August 2016

A thesis submitted to McGill University in partial fulfilment
of the requirements of the degree of Master of Science

Abstract

Revealed preference techniques are used to test whether a data set is compatible with rational behaviour. They are also incorporated as constraints in mechanism design to encourage truthful behaviour. In auctions, revealed preference bidding constraints are very demanding. We explain first how additively approximate rationality can be used to create relaxed revealed preference constraints in an auction, and show how combinatorial methods can be used to implement these relaxed constraints. Worst/best-case welfare guarantees that result from the use of such mechanisms can be quantified via the minimum/maximum virtual valuation function. We then go on to consider a combinatorial measure of approximate rationality, and show that computing the degree of rationality is NP-hard, except in a 2-commodity market. In showing that it is efficiently computable in 2-commodity markets, we introduce a class of perfect graphs, which we believe to be new. When the market has at least three commodities, we show that the problem is NP-complete by a reduction from 3-SAT. To complete this reduction, we introduce the class of *oriented-disk graphs.*

Abrégé

On utilise les techniques de la préférence révélée pour déterminer si une collection de données peut être modélisée par les choix d'un consommateur rationnel. Ces techniques sont aussi parfois mises en place pour encourager le comportement véridique. Chez les enchères, les contraintes de demande à base de préférence révélée sont plutôt exigeantes. D'abord, nous expliquons comment le comportement rationnel arithméthiquement approximatif peut être utilisé pour créer une contrainte de demande moins exigeante, et produisons une méthode combinatoire pour l'implémenter. De plus, les garanties minimales et maximales de bien-être du consommateur peuvent être quantifiées en calculant les utilités virtuelles minimalement et maximalement réalisables. Par la suite, nous considérons le comportement rationnel combinatoirement approximatif, et déterminons que calculer le niveau d'approximation est un problème NP-difficile, sauf lorsqu'il y a seulement deux items au marché. Pour démontrer que le problème peut être résolu efficacement le cas échéant, nous introduisons une classe de graphes parfaits, que l'on soupçonne être inconnu à date. Lorsqu'il-y-a au moins trois items au marché, nous démontrons l'NP-complétude du problème en réduisant à partir de 3-SAT. Cette réduction nécessite l'introduction de la classe des *graphes à disques orientés*.

Acknowledgements

I would like to thank my supervisor Adrian Vetta for all his involvement. With his help, I was able to write and publish two conference submissions in under a year's time; I will always be grateful to him for encouraging me to publish as much work as possible. In fact, those two articles form the body of this thesis. I would also like to thank my various office-mates throughout my degree, and my friends outside of school, who were instrumental in keeping me sane throughout the process of writing those articles and this thesis.

I would also like to thank Sergey Norin for suggesting a simplification to the proof of Theorem 3.2 in his external report. This proof has been included as Theorem 3.5.

This thesis was written while I was receiving financial support from the Department of Mathematics and Statistics, via my supervisor Adrian Vetta.

Contribution of Authors

The body of this thesis is the result of joint work with prof. Adrian Vetta. Chapter 2 is an edited version of "The Combinatorial World (of Auctions) According to GARP," published in the proceedings of the *8th International Symposium on Algorithmic Game Theory* (SAGT '15); and Chapter 3 is an edited version of "Testing Consumer Rationality Using Perfect Graphs and Oriented Disks," published in the proceedings of the *11th Conference on Web and Internet Economics* (WINE '15). Both papers were authored by myself and Adrian Vetta.

Contents

Chapter 1

Introduction

Historically, consumer theory has sought to model the behaviour of an economically rational consumer. Beginning with Samuelson in 1938 [34], efforts have been made to characterise the behaviour of a rational agent via *revealed preferences*: the notion that a consumer's choices reveal information about her underlying preferences. This gave rise to the *axioms of revealed preference* outlined in Section 1.1, and culminated with Afriat's Theorem [1], which exactly characterises rational behaviour. See also [41] for more details. The axioms of revealed preference formalise the observation that the actions of a rational agent must be consistent with her historical behaviour in the market.

These axioms have been used in mechanism design to encourage rational behaviour in agents. Consider, for example, combinatorial auctions *i.e.* simultaneous auctions of a collection of indivisible items. These auctions have become a popular method of allocating many resources which are related to one another. Such auctions, however, are susceptible to manipulation by the bidders: they may underbid in early rounds to keep prices low and hide their true interests [6]. As a result of such behaviour, many *activity rules* have been developed; these rules restrict the bidders' behaviour so as to make dishonest behaviour either impossible or non-profitable. To counter the problem of underbidding for example, a monotonicity requirement was first implemented: bidders wishing to bid for large quantities at the end of the auction must, essentially, bid for large quantities in the early rounds. However, such a rule on its own may hinder some bidders from bidding truthfully. Thus, in 2006, Ausubel, Crampton and Milgrom [6] introduced a rule which permitted agents to violate the monotonicity requirement so long as their behaviour was consistent with past behaviour. They determined whether an agent was being consistent via the *weak axiom of revealed prefence,* or WARP. (See Section 1.1) This, however, only tests for consistency in pairs of data points. In Chapter 2, we extend their derivation to the *generalised axiom* (GARP), and provide a graphical framework which allows us to efficiently verify that a bidder is behaving consistently. We then derive a related test to determine whether a bidder is behaving almost-consistently, in the sense that the bids they have placed are within some additive ε of their optimal bid. This test reduces to computing the *minimum mean cycle* (MMC) on the same graph. We go on to prove that, given the value of the MMC, there exists a unique minimum (and maximum) estimate for the agent's valuation among all valid estimates, and provide a polytime algorithm to compute it. Finally, we note that the tests for rationality and ε-approximate rationality can be implemented as bidding rules, since they are poly-time computable. We then quantify the loss in the guarantee of rationality if we bound the number of bids required to show a violation.

Various measures have been constructed to determine not only whether an agent is being rational, but more specifically, the degree to which she is irrational. See [40] for an overview. We outline a few of these in concluding Chapter 2, including a measure of rationality based on the number of bids needed to be ignored for the remaining bids to be consistent. This leads us to Chapter 3, where we consider such a measure in a broader context. Since rationality in standard consumer theory is equivalent to acyclicity in preference (see Section 1.1) the above problem can be modelled as the *minimum directed feedback vertex set* (MDFVS) problem on directed preference graphs. It is known that, in full generality, any directed graph is a feasible preference graph. However, if we restrict the number of items in the market, *i.e.* the dimension of the data points, the class of constructible graphs is also restricted. We show that in a 2-commodity market, the problem can be reduced to solving the *minimum vertex cover* (MVC) problem on an auxiliary graph, which we then show is perfect. Hence, by the famous results of Grötschel, Lovász and Schrijver [20, 21], the problem can be solved in polytime. (In fact, this class of auxiliary graphs appears to be a previously unknown class of perfect graphs.) We then show that in a 3-commodity market, the problem becomes NP-hard, via a reduction from 3-SAT. To do this, we introduce the class of *oriented-disk graphs* which we show is a subset of preference graphs feasible in 3-commodity markets.

§1.1 The Axioms of Revealed Preference

In this section we provide a brief overview of the development of the *axioms of revealed preference*. We consider, for this section and for Chapter 3, a rational agent purchasing a bundle in a market of n different divisible resources. Any bundle of goods in this market can be viewed as a vector in $\mathbf{R}^n$, where each component is the quantity of some good in the bundle. (Here $\mathbf{R}$ denotes the non-negative reals.) Our agent, according to standard consumer theory, has a valuation function $v : \mathbf{R}^n \to \mathbf{R}$ and a budget $B \in \mathbf{R}$. If in this market, goods are priced linearly, we can view prices as a vector $\boldsymbol{p} \in \mathbf{R}^n$; the price of a bundle $\boldsymbol{x} \in \mathbf{R}^n$ is then the inner product $\boldsymbol{p} \cdot \boldsymbol{x}$. In this setting, the agent is said to be *rational* if their demanded bundle is the most valuable bundle subject to the budget constraint, *i.e.*

$$\mathit{demand}(v; \boldsymbol{p}) \;=\; \operatorname*{arg\,max}_{x : \boldsymbol{p} \cdot \boldsymbol{x} \leq B} v(\boldsymbol{x}) \;. \tag{1.1.1}$$

Now, suppose we have a collection of observed *consumer data* $\{(\boldsymbol{p}_1, \boldsymbol{x}_1), (\boldsymbol{p}_2, \boldsymbol{x}_2), \ldots, (\boldsymbol{p}_m, \boldsymbol{x}_m)\}$. Each pair $(\boldsymbol{p}_i, \boldsymbol{x}_i)$ denotes the fact that the consumer purchased the bundle of goods $\boldsymbol{x}_i \in \mathbf{R}^n$ when the prices were $\boldsymbol{p}_i \in \mathbf{R}^n$. If we then have $\boldsymbol{p}_i \cdot \boldsymbol{x}_i \leq \boldsymbol{p}_i \cdot \boldsymbol{x}_j$, we have that $\boldsymbol{x}_j$ is affordable at prices $\boldsymbol{p}_i$, and therefore, by the choice of $\boldsymbol{x}_i$, the agent has *revealed* that $\boldsymbol{x}_i$ is *preferred* to $\boldsymbol{x}_j$ (when both are affordable). We say $\boldsymbol{x}_i$ is *directly revealed preferred* to $\boldsymbol{x}_j$ and denote this $\boldsymbol{x}_i \succeq \boldsymbol{x}_j$. Furthermore, suppose we observe that $\boldsymbol{x}_i \succeq \boldsymbol{x}_j$ and that $\boldsymbol{x}_j \succeq \boldsymbol{x}_k$. Then, by transitivity of preference, we say $\boldsymbol{x}_i$ is *indirectly revealed preferred* to $\boldsymbol{x}_k$.

It is important to note here that the *demand* function in Equation 1.1.1 may be *single-valued* or *multi-valued*, *i.e.* there may be always one unique demanded item, or a set of demanded items. If it is single-valued, purchasing $\boldsymbol{x}_i$ when $\boldsymbol{x}_j$ was affordable reveals that $\boldsymbol{x}_i$ is strictly preferable to $\boldsymbol{x}_j$, denoted $\boldsymbol{x}_i \succ \boldsymbol{x}_j$. However, if it is multi-valued, it is possible that the agent was indifferent between the bundles, and hence we may only conclude $\boldsymbol{x}_i \succeq \boldsymbol{x}_j$.

In 1938, Samuelson [34] noted that, if an agent reveals both $\boldsymbol{x}_i \succ \boldsymbol{x}_j$ and $\boldsymbol{x}_j \succ \boldsymbol{x}_i$, then the agent is clearly behaving inconsistently. From this, we conclude that a rational agent may not behave as such. This was later termed the *weak axiom of revealed preference* (WARP) for the behaviour of a rational agent.

The Weak Axiom of Revealed Preference. If $\boldsymbol{x}_1 \succ \boldsymbol{x}_2$, then $\boldsymbol{x}_2 \not\succ \boldsymbol{x}_1$.

This is clearly a necessary condition for rationality. However Samuelson claimed that this was also a sufficient condition and proved this in the trivial case of $m = 2$, where m is the number of data points. Houtthakker [26] later showed that this was in fact only true when $m = 2$, and instead considered the consistency of *indirectly* revealed preferences. The resulting axiom was termed the *strong axiom of revealed preference* (SARP).

The Strong Axiom of Revealed Preference. If $\boldsymbol{x}_1 \succ \boldsymbol{x}_2$, $\boldsymbol{x}_2 \succ \boldsymbol{x}_3$, $\ldots$, $\boldsymbol{x}_{k-2} \succ \boldsymbol{x}_{k-1}$ and $\boldsymbol{x}_{k-1} \succ \boldsymbol{x}_k$ then $\boldsymbol{x}_k \not\succ \boldsymbol{x}_1$.

Houtthakker proved that SARP was necessary and sufficient for rationality if demand is always single-valued. One can expect, though, that agents have multi-valued demand; however, the two previous axioms are defined in terms of strict preference, and as such are not suited for this setting. In 1967, Afriat [1] constructed the *generalised axiom* (GARP) below, and showed that it is a necessary and sufficient condition for rationality in this more general setting. This result is now known as *Afriat's Theorem.*

The Generalised Axiom of Revealed Preference. If $\boldsymbol{x}_1 \succeq \boldsymbol{x}_2$, $\boldsymbol{x}_2 \succeq \boldsymbol{x}_3$, $\ldots$, $\boldsymbol{x}_{k-1} \succeq \boldsymbol{x}_k$ and $\boldsymbol{x}_k \succeq \boldsymbol{x}_1$, then $\boldsymbol{x}_1 \sim \boldsymbol{x}_k$. (Equivalently, $\boldsymbol{x}_k \not\succ \boldsymbol{x}_1$.)

Afriat showed that this was a sufficient condition by giving a method to construct monotonic, concave, piecewise-linear utility functions from any data set satisfying GARP. It is easy to check that this is a necessary condition.

Afriat's Theorem is an important result: it allows us to efficiently test whether a consumer is behaving consistently by simply checking for cycles in the preference data. In the rest of this thesis, we explore derivations of similar results, and modifications of the axiom which test not only the agent's consistency, but also quantifies the degree to which they are consistent.

For the purposes of relaxations in Chapter 2, we introduce here the *k-th Axiom of Revealed Preference*:

The k-th Axiom of Revealed Preference. For any $\kappa \leq k+1$,
if $\boldsymbol{x}_1 \succeq \boldsymbol{x}_2$, $\boldsymbol{x}_2 \succeq \boldsymbol{x}_3$, $\ldots$, $\boldsymbol{x}_{\kappa-1} \succeq \boldsymbol{x}_\kappa$ and $\boldsymbol{x}_\kappa \succeq \boldsymbol{x}_1$, then $\boldsymbol{x}_\kappa \not\succ \boldsymbol{x}_1$.

We note that this is very similar to GARP. In fact, this is exactly GARP if k is taken to be arbitrarily large. Furthermore, we have also that, for $k = 1$, this is a multivalued-demand version of WARP. Thus, the value of k parametrises the entire spectrum of axioms of revealed preference between the weak and strong axiom.

Chapter 2

An Additive Measure of Rationality in the Quasilinear Setting: Tractable Bidding Rules on a Versatile Graph

As discussed in Chapter 1, revealed preference has become an important tool in auction design, since bidding rules based on WARP have been implemented. These rules are now standard in the combinatorial clock auction, one of the two prominent auction mechanisms used to sell bandwidth. In part, the WARP-based bidding rules have proved successful because they are extremely difficult to game [10]. However, Harsha et al. [24] examine GARP-based bidding rules, and Ausubel and Baranov [5] advocate incorporating such constraints into bandwidth auctions. Based upon Afriat's theorem, these GARP-based rules imply that there always exists a utility function that is compatible with the bidding history. This gives the desirable property that a bidder in an auction will always have at least one feasible bid – a property that cannot be guaranteed under WARP.

In this chapter, we show how a graphical viewpoint of revealed preference can be used to obtain a virtual valuation function that best fits the data set. Specifically, we show in Section 2.3 that an individually rational virtual valuation function can be obtained such that its additive deviation from rationality is exactly the *minimum mean length* of a cycle in a bidding graph. This additive guarantee cannot be improved upon. Furthermore, we show there exists a unique *minimum* valuation function from amongst all individually rational virtual valuation functions that optimally fit the data. Similarly, given a set of upper bound constraints, we show how to find the unique *maximum* virtual valuation that optimally fits the data, if it exists.

Imposing revealed preference bidding rules can be harsh. Indeed, Cramton [10] states that "there are good reasons to simplify and somewhat weaken the revealed preference rule". These reasons include complexity issues, common value uncertainty, the complication of budget constraints, and the fact that a bidder's assessment of her valuation function often *changes* as the auction progresses! The concept of approximate rationality, however, naturally induces a relaxed form of revealed preference rules. We examine such relaxed bidding rules in Section 2.5, show how they can be implemented combinatorially, and show how to construct the minimal and maximal valuation functions which fit the data, which may be useful for quantifying worst-case and best-case welfare guarantees.

§2.1 Revealed Preference in Combinatorial Auctions

As discussed, a major application of revealed preference in mechanism design concerns combinatorial auctions. Here, there are some important distinctions from the standard revealed preference model presented in Section 1.1[1]. First, consumers are assumed to have quasilinear utility functions that are linear in money. Thus, they seek to maximise profit. Second, the standard assumption is that bidders have *no* budgetary constraints. For example, if profitable opportunities arise that require large investments then these can be obtained. (This

[1] However, as explained below, this slightly different model can be seen as a special case of the more general model discussed in the first chapter.

assumption is slightly unrealistic; Harsha et al. [24] show how to implement a budgeted revealed preference model for combinatorial auctions; see also Section 2.6).

Third, the observations $(\boldsymbol{p}_t, \boldsymbol{x}_t)$, for each $1 \leq t \leq T$, are typically not purchases but are bids made over a collection of auction rounds. When offered a set of prices at time t the consumer bids for bundle $\boldsymbol{x}_t$. In such auctions, the market typically consists of a collection of indivisible items, rather than a commodity market (*i.e.* bundles are represented by 0-1 vectors.) However, we will still be denoting bundles as general vectors, as most of the following results still hold in the general setting.

So what would a model of revealed preference be in this combinatorial auction setting? Suppose that at time t we select bundle $\boldsymbol{x}_t$ and that at an earlier time s we selected bundle $\boldsymbol{x}_s$. Assuming a quasi-linear utility function and no budget constraint, we have revealed:

$$v(\boldsymbol{x}_t) - \boldsymbol{p}_t \cdot \boldsymbol{x}_t \geq v(\boldsymbol{x}_s) - \boldsymbol{p}_t \cdot \boldsymbol{x}_s \tag{2.1.1}$$

$$v(\boldsymbol{x}_s) - \boldsymbol{p}_s \cdot \boldsymbol{x}_s \geq v(\boldsymbol{x}_t) - \boldsymbol{p}_s \cdot \boldsymbol{x}_t \tag{2.1.2}$$

Summing Inequalities (2.1.1) and (2.1.2) and rearranging gives

$$(\boldsymbol{p}_t - \boldsymbol{p}_s) \cdot \boldsymbol{x}_s \geq (\boldsymbol{p}_t - \boldsymbol{p}_s) \cdot \boldsymbol{x}_t \tag{2.1.3}$$

This is the revealed preference condition for combinatorial auctions proposed as a bidding activity rule by Ausubel, Crampton and Milgrom [6]. The activity rule simply states that, between time s and time t, the price of bundle $\boldsymbol{x}_t$ must have risen by at least as much as the price of $\boldsymbol{x}_s$. If condition (2.1.3) is not satisfied then the auction mechanism will not allow the later bid to be made.

Observe that the bidding rule (2.1.3) was derived directly from the assumption of utility maximisation. This unbudgeted revealed preference auction model can, though, also be viewed within the framework of the standard budgeted model of revealed preference. To do this, we assume the bidder has an arbitrarily large budget B. In particular, prices will never be so high that she cannot afford to buy every item. Second, to model quasilinear utility functions, we treat money as a good. Specifically, given a bundle of items $\boldsymbol{x} = (x_1, \ldots, x_n)$ and an amount x_0 of money we denote by $\hat{\boldsymbol{x}} = (x_0, x_1, \ldots, x_n)$ the concatenation of x_0 and $\boldsymbol{x}$. If $\boldsymbol{p} = (p_1, \ldots, p_n)$ is the price vector for the the non-monetary items, then $\hat{\boldsymbol{p}} = (1, p_1, \ldots, p_n)$ gives the prices of all items including money.

In this $n+1$ dimensional setting, let us select bundle $\hat{\boldsymbol{x}}_t$ at time t. As the budget B is arbitrarily large, we can certainly afford the bundle $\boldsymbol{x}_s$ at this time. But we may not be able to afford bundle $\hat{\boldsymbol{x}}_s$, as then we must also pay for the monetary component at a cost of $B - \boldsymbol{p}_s \cdot \boldsymbol{x}_s$. However, we can afford the bundle $\boldsymbol{x}_s$ plus an amount $B - \boldsymbol{p}_t \cdot \boldsymbol{x}_s$ of money. Applying revealed preference to $\{\hat{\boldsymbol{x}}, \hat{\boldsymbol{p}}\}$, we have revealed that $\hat{\boldsymbol{x}}_t = (B - \boldsymbol{p}_t \cdot \boldsymbol{x}_t, \boldsymbol{x}_t) \succeq (B - \boldsymbol{p}_t \cdot \boldsymbol{x}_s, \boldsymbol{x}_s)$. Hence, by quasilinearity, subtracting the monetary component from both sides, we have,

$$(0, \boldsymbol{x}_t) \;\succeq\; ((B - \boldsymbol{p}_t \cdot \boldsymbol{x}_s) - (B - \boldsymbol{p}_t \cdot \boldsymbol{x}_t), \boldsymbol{x}_s) \;=\; (\boldsymbol{p}_t \cdot \boldsymbol{x}_t - \boldsymbol{p}_t \cdot \boldsymbol{x}_s, \boldsymbol{x}_s) \;\; .$$

Equivalently,

$$v(\boldsymbol{x}_t) \geq v(\boldsymbol{x}_s) + \boldsymbol{p}_t \cdot \boldsymbol{x}_t - \boldsymbol{p}_t \cdot \boldsymbol{x}_s \;\; . \tag{2.1.4}$$

But Inequality (2.1.4) is equivalent to Inequality (2.1.1). Inequality (2.1.2) follows symmetrically, and together these give the revealed preference bidding rule (2.1.3). Note that this bidding rule is derived via the direct comparison of two bundles.

We can now extend this bidding rule to incorporate indirect comparisons in a similar fashion to the extension from WARP to SARP via transitivity. This produces a GARP-based bidding rule. Namely, suppose we bid for the money-less bundle $\boldsymbol{x}_i$ at time t_i, for all $0 \leq i \leq k$, where $1 \leq t_i \leq T$. Thus we have revealed that

$$\begin{aligned}(0, \boldsymbol{x}_i) &\succeq ((B - \boldsymbol{p}_i \cdot \boldsymbol{x}_{i+1}) - (B - \boldsymbol{p}_i \cdot \boldsymbol{x}_i), \boldsymbol{x}_{i+1}) \\ &= (\boldsymbol{p}_i \cdot \boldsymbol{x}_i - \boldsymbol{p}_i \cdot \boldsymbol{x}_{i+1}, \boldsymbol{x}_{i+1})\end{aligned}$$

This induces the inequality

$$v(\boldsymbol{x}_i) - \boldsymbol{p}_i \cdot \boldsymbol{x}_i \geq v(\boldsymbol{x}_{i+1}) - \boldsymbol{p}_i \cdot \boldsymbol{x}_{i+1} \ . \tag{2.1.5}$$

Summing (2.1.5) over all i, we obtain

$$\sum_{i=0}^{k} (v(\boldsymbol{x}_i) - \boldsymbol{p}_i \cdot \boldsymbol{x}_i) \ \geq \ \sum_{i=0}^{k} (v(\boldsymbol{x}_{i+1}) - \boldsymbol{p}_i \cdot \boldsymbol{x}_{i+1}) \ \ ,$$

where the sum in the subscripts are taken modulo k. Rearranging now gives the combinatorial auction KARP-based bidding activity rule:

$$(\boldsymbol{p}_k - \boldsymbol{p}_0) \cdot \boldsymbol{x}_0 \ \geq \ \sum_{i=1}^{k} (\boldsymbol{p}_i - \boldsymbol{p}_{i-1}) \cdot \boldsymbol{x}_i \ . \tag{2.1.6}$$

For k arbitrarily large, this gives the GARP-based bidding rule. In order to qualitatively analyze the consequences of imposing KARP-based activity rules, it is informative to now provide a graphical interpretation of the these rules.

§2.2 A Graphical View of Revealed Preference

Given the bidder data $\{(\boldsymbol{p}_t, \boldsymbol{x}_t) : 1 \le t \le T\}$, we create a directed graph $G = (V, A)$, called the *bidding graph*, to which we will assign arc lengths ℓ. There is a vertex in V for each possible bundle – that is, there are 2^n bundles in an n-item auction For each observed bid $\boldsymbol{x}_t$, $1 \le t \le T$, there is an arc $(\boldsymbol{x}_t, \boldsymbol{y})$ for each bundle $\boldsymbol{y} \in V$. In order to define the length $\ell_{\boldsymbol{x}_t,\boldsymbol{y}}$ of an arc $(\boldsymbol{x}_t, \boldsymbol{y})$, note that Inequality (2.1.1) applied to $\boldsymbol{x}_s = \boldsymbol{y}$ gives

$$v(\boldsymbol{y}) \;\le\; v(\boldsymbol{x}_t) + \boldsymbol{p}_t \cdot (\boldsymbol{y} - \boldsymbol{x}_t) \;\;,$$

otherwise we would prefer bundle $\boldsymbol{y}$ at time t. For the arc length, we would like to simply set $\ell_{\boldsymbol{x}_t,\boldsymbol{y}} = \boldsymbol{p}_t \cdot (\boldsymbol{y} - \boldsymbol{x}_t)$. Observe, however, that the bundle $\boldsymbol{x}_t$ may be chosen in more than one time period. That is, possibly $\boldsymbol{x}_t = \boldsymbol{x}_{t'}$ for some $t \ne t'$. Therefore the bidding graph is, in fact, a multigraph. It suffices, though, to represent only the most stringent constraints imposed by the bidding behaviour. Thus, we obtain a simple graph by setting

$$\ell_{\boldsymbol{x}_t,\boldsymbol{y}} \;=\; \min_{t'} \{\boldsymbol{p}_{t'} \cdot (\boldsymbol{y} - \boldsymbol{x}_t) : \boldsymbol{x}_{t'} = \boldsymbol{x}_t\} \;\;.$$

Now the WARP-based bidding rule (2.1.3) of Ausubel et al. [6] is equivalent to

$$(\boldsymbol{p}_t - \boldsymbol{p}_s) \cdot \boldsymbol{x}_s - (\boldsymbol{p}_t - \boldsymbol{p}_s) \cdot \boldsymbol{x}_t \ge 0 \;\;.$$

However,

$$\begin{aligned}
&\ell_{\boldsymbol{x}_s,\boldsymbol{x}_t} + \ell_{\boldsymbol{x}_t,\boldsymbol{x}_s} \\
&\quad = \min_{s'} \{\boldsymbol{p}_{s'} \cdot (\boldsymbol{x}_t - \boldsymbol{x}_s) : \boldsymbol{x}_{s'} = \boldsymbol{x}_s\} + \min_{t'} \{\boldsymbol{p}_{t'} \cdot (\boldsymbol{x}_s - \boldsymbol{x}_t) : \boldsymbol{x}_{t'} = \boldsymbol{x}_t\} \\
&\quad \le \boldsymbol{p}_s \cdot (\boldsymbol{x}_t - \boldsymbol{x}_s) + \boldsymbol{p}_t \cdot (\boldsymbol{x}_s - \boldsymbol{x}_t) \\
&\quad = (\boldsymbol{p}_t - \boldsymbol{p}_s) \cdot \boldsymbol{x}_s - (\boldsymbol{p}_t - \boldsymbol{p}_s) \cdot \boldsymbol{x}_t \;\;.
\end{aligned}$$

It is then easy to see that the bidding constraint (2.1.3) is violated if and only if the bidding graph contains no negative digons (cycles of length two). Furthermore, we can interpret KARP and GARP is a similar fashion. Hence, the k-th axiom of revealed preference is equivalent to requiring that the bidding graph not contain any negative cycles of cardinality at most $k+1$, and GARP is equivalent to requiring no negative cycles at all. Thus, we can formalize the preference axioms in terms of the lengths of negative cycles in a directed graph. We remark that a cyclic view of revealed preference is briefly outlined by Vohra [43]. For us, this cyclic formulation has important consequences in testing for the extent of bidding deviations from the axioms. We will quantify this exactly in Section 2.3. Before doing so, though, we remark that the focus on cycles also has important computational consequences.

First, recall that the bidding graph G contains an exponential number of vertices, one for every subset of the items. Of course, it is not practical to work with such a graph. Observe, however, that a bundle $\boldsymbol{y} \notin \{\boldsymbol{x}_1, \boldsymbol{x}_2 \ldots, \boldsymbol{x}_T\}$ has zero out-degree in G. Consequently, $\boldsymbol{y}$ cannot be contained in any cycle. Thus, it will suffice to consider only the subgraph induced by the bids $\{\boldsymbol{x}_1, \boldsymbol{x}_2 \ldots, \boldsymbol{x}_T\}$. In a combinatorial auction there is typically one bid per time period and the number of periods is quite small.[2] Hence, the induced subgraph of the bidding graph that we actually need is of a very manageable size.

Second, one way to implement a bidding rule is via a mathematical program; see, for example, Harsha et al. [24]. The cyclic interpretation of a bidding rule has two major advantages: we can test the rule very quickly by searching for negative cycles in a graph. For example, we can test for negative cycles of length at most $k+1$ either by fast matrix multiplication or directly by looking for shortest paths of length k using the Bellman-Ford algorithm in

[2] For example, in a bandwidth auction there are at most a few hundred rounds.

$O(T^3)$ time. Another major advantage is that a bidder can interpret the consequence of a prospective new bid dynamically by consideration of the bidding graph. This is extremely important in practice. In contrast, bidding rules that require using an optimization solver as a black-box are very opaque to bidders.

§2.3 Minimum Mean Cycles and Approximate Virtual Valuation Functions

For combinatorial auctions, Afriat's result that GARP is necessary and sufficient for rationalisability can be reformulated as:

Theorem 2.1. *A valuation function which rationalises bidding behaviour exists if and only if the bidding graph has no negative cycle.*

This is a simple corollary of Theorem 2.2 below; see also [43]. From an economic perspective, however, what is most important is not whether agents are perfectly rational but "whether optimization is a reasonable way to describe some behavior" [40].[3] It is then important to study the consequences of approximately rational behaviour, see, for example, Akerlof and Yellen [3]. First, though, is it possible to quantify the degree to which agents are rational? Gross [19] examines assorted methods to test the degree of rationality. Notable amongst them is the *Afriat Efficiency Index* [1, 40]. Here the condition required to imply a preference is strengthened multiplicatively. Specifically, $\boldsymbol{x}_t \succeq \boldsymbol{y}$ only if $\boldsymbol{p}_t \cdot \boldsymbol{y} \leq \lambda \cdot \boldsymbol{p}_t \cdot \boldsymbol{x}_t$ where $\lambda < 1$. We examine this index with respect to the bidding graph in Section 2.6. For combinatorial auctions, a variant of this constraint was examined experimentally by Harsha et al. [24].

[3] Indeed, several schools of thought in the field of bounded rationality argue that people utilize simple (but often effective) heuristics rather than attempt to optimize; see, for example, [18].

Here we show how to quantify exactly the degree of rationality present in the data via a parameter of the bidding graph. Moreover, we are able to go beyond multiplicative guarantees and obtain stronger additive bounds. To wit, we say that $\hat{v}$ is an ϵ-*approximate virtual valuation function* if, for all t and for any bundle $\boldsymbol{y}$,

$$\hat{v}(\boldsymbol{x}_t) - \boldsymbol{p}_t \cdot \boldsymbol{x}_t \geq \hat{v}(\boldsymbol{y}) - \boldsymbol{p}_t \cdot \boldsymbol{y} - \epsilon \ .$$

Note that if $\epsilon = 0$, then the bidder is optimizing with respect to a virtual valuation function, *i.e.* is rational. We remark that the term *virtual* reflects the fact that $\hat{v}$ need not be the real valuation function (if one exists) of the bidder, but if it is then the bidding is termed *truthful*.

We now examine exactly when a bidding strategy is approximately rational. It turns out that the key to understanding approximate deviations from rationality is the *minimum mean cycle* in the bidding graph. Given a cycle C in G, its mean length is

$$\mu(C) = \frac{\sum_{a \in C} \ell_a}{|C|} \ .$$

We denote by $\mu(G) = \min_C \mu(C)$ the *minimum mean length* of a cycle in G, and we say that C^* is a *minimum mean cycle* if $C^* \in \operatorname{argmin}_C \mu(C)$. We can find a minimum mean cycle in polynomial time using the classical techniques of Karp [29].

Theorem 2.2. *An ϵ-approximate valuation function which (approximately) rationalises bidding behaviour exists if and only if the bidding graph has minimum mean cycle $\mu(G) \geq -\epsilon$.*

Proof. From the bidding graph G we create an auxiliary directed graph $\hat{G} = (\hat{V}, \hat{A})$ with vertex set $\hat{V} = \{\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_T\}$. The arc set is complete with arc lengths

$$\hat{\ell}_{\boldsymbol{x}_s,\boldsymbol{x}_t} = \ell_{\boldsymbol{x}_s,\boldsymbol{x}_t} - \mu(G) \ .$$

Observe that, by construction, every cycle in $\hat{G}$ is of non-negative length. It follows that we may obtain shortest path distances $\hat{d}$ from any arbitrary root vertex r. Thus, for any arc $(\boldsymbol{x}_t, \boldsymbol{y})$, we have

$$\begin{aligned} \hat{d}(\boldsymbol{y}) &\leq \hat{d}(\boldsymbol{x}_t) + \hat{\ell}_{\boldsymbol{x}_t,\boldsymbol{y}} \\ &= \hat{d}(\boldsymbol{x}_t) + \ell_{\boldsymbol{x}_t,\boldsymbol{y}} - \mu(G) \\ &\leq \hat{d}(\boldsymbol{x}_t) + \boldsymbol{p}_t \cdot (\boldsymbol{y} - \boldsymbol{x}_t) - \mu(G) \ . \end{aligned}$$

So, if we set $\hat{v}(\boldsymbol{x}) = \hat{d}(\boldsymbol{x})$, for each $\boldsymbol{x}$, then

$$\hat{v}(\boldsymbol{x}_t) - \boldsymbol{p}_t \cdot \boldsymbol{x}_t \geq \hat{v}(\boldsymbol{y}) - \boldsymbol{p}_t \cdot \boldsymbol{y} + \mu(G) \ .$$

for all t. Therefore, by definition of ϵ-approximate bidding, we have that $\hat{v}$ is a $(-\mu)$-approximate virtual valuation function.

Conversely, let $\hat{v}$ be an ϵ-approximate virtual valuation function which rationalises the graph, and take some cycle C of minimum mean length in the bidding graph. Suppose for a contradiction that $\mu(C) < -\epsilon$. By ϵ-approximability, we have

$$\hat{v}(\boldsymbol{x}_s) - \boldsymbol{p}_s \cdot \boldsymbol{x}_s \geq \hat{v}(\boldsymbol{x}_t) - \boldsymbol{p}_s \cdot \boldsymbol{x}_t - \epsilon \ .$$

But $\ell_{\boldsymbol{x}_s,\boldsymbol{x}_t} \geq \boldsymbol{p}_s \cdot (\boldsymbol{x}_t - \boldsymbol{x}_s)$. Therefore $\ell_{\boldsymbol{x}_s\boldsymbol{x}_t} \geq \hat{v}(\boldsymbol{x}_t) - \hat{v}(\boldsymbol{x}_s) - \epsilon$. Summing over every arc in the cycle we obtain

$$\ell(C) \;=\; \sum_{(\boldsymbol{x},\boldsymbol{y})\in C} \ell_{\boldsymbol{xy}} \;\geq\; \sum_{(\boldsymbol{x},\boldsymbol{y})\in C} (\hat{v}(\boldsymbol{y}) - \hat{v}(\boldsymbol{x}) - \epsilon) \;=\; -|C| \cdot \epsilon \;\; .$$

Thus $\mu(C) \geq -\epsilon$, giving the desired contradiction. □

Recall that, the bidding behaviour is irrational only if $\mu(G)$ is strictly negative. We emphasize that Theorem 2.2 applies even when $\mu(G)$ is positive, but in this case, we have an ϵ-approximate virtual valuation function where ϵ is negative! What does this mean? Well, setting $\delta = -\epsilon$, we then have, for all t and for any bundle $\boldsymbol{y}$, that $\hat{v}(\boldsymbol{x}_t) - \boldsymbol{p}_t \cdot \boldsymbol{x}_t \;\geq\; \hat{v}(\boldsymbol{y}) - \boldsymbol{p}_t \cdot \boldsymbol{y} + \delta$. Thus, $\boldsymbol{x}_t$ is not just the best choice, but it provides at least an extra δ units of utility over any other bundle. Thus, the larger δ is, the greater our degree of confidence in the revealed preference-ordering and valuation.

§2.4 Minimum & Maximum Individually Rational Virtual Valuation Functions

Theorem 2.2 shows how to obtain a virtual valuation function with the best possible additive approximation guarantee: any valuation rationalising the bidding graph G must allow for an additive approximation of at least $-\mu(G)$. However, there is a problem. Such a valuation function may not actually be compatible with the data; specifically, it may not be individually rational. For *individual rationality*, we require, for each time t, that $\hat{v}(\boldsymbol{x}_t) - \boldsymbol{p}_t \cdot \boldsymbol{x}_t \geq 0$. But individually rationality is (almost certainly) violated for the the root node r since we have $\hat{v}(\boldsymbol{x}_r) = 0$.

It is possible to obtain an individually rational, approximate, virtual valuation function simply by taking the $\hat{v}$ from Theorem 2.2 and adding a huge constant to value of each package. This operation, of course, is entirely unnatural and the resulting valuation function is of little practical value.

We say that $v()$ is the *minimum individually rational, ϵ-approximate virtual valuation function* if $v(\boldsymbol{x}_t) \leq \omega(\boldsymbol{x}_t)$ for each $1 \leq t \leq T$, for any other individually rational, ϵ-approximate virtual valuation function $\omega()$. This leads to the questions: (i) Does such a valuation function exist? and (ii) Can it be obtained efficiently? The answer to both these questions is *yes*.

Theorem 2.3. *The minimum individually rational, μ-approximate virtual valuation function exists and can be found in polynomial time.*

Proof. Let $\hat{G}$ be as in Theorem 2.2. We create an auxiliary directed graph H from $\hat{G}$ by adding a sink vertex $\boldsymbol{z}$. We add an arc $(\boldsymbol{x}_t, \boldsymbol{z})$ of length $-\boldsymbol{p}_t \cdot \boldsymbol{x}_t$, for each $1 \leq t \leq T$, allowing for repeated arcs. Because $\hat{G}$ contains no negative cycle, neither does H. Therefore, there exist shortest path distances in H. Denote by $\hat{d}()$ the shortest path distance from vertex $\boldsymbol{x}_t$ to $\boldsymbol{z}$ in H. We claim that setting $v(\boldsymbol{x}_t) = -\hat{d}(\boldsymbol{x}_t)$ gives the minimum individually rational, μ-approximate virtual valuation function.

To begin, let's verify that $v()$ is an individually rational, μ-approximate virtual valuation function. First, we require that $v()$ is individually rational. Now the direct path consisting of the arc $(\boldsymbol{x}_t, \boldsymbol{z})$ is at least as long as the shortest path from $\boldsymbol{x}_t$ to $\boldsymbol{z}$. Thus, $-\boldsymbol{p}_t \cdot \boldsymbol{x}_t \geq \hat{d}(\boldsymbol{x}_t)$. Individual rationality then follows as $v(\boldsymbol{x}_t) = -\hat{d}(\boldsymbol{x}_t) \geq \boldsymbol{p}_t \cdot \boldsymbol{x}_t$.

Second we need to show that $v()$ is μ-approximate. Consider a pair $\{\boldsymbol{x}_s, \boldsymbol{x}_t\}$. The shortest path conditions imply that

$$-v(\boldsymbol{x}_s) \;=\; \hat{d}(\boldsymbol{x}_s) \;\leq\; \hat{\ell}_{st} + \hat{d}(\boldsymbol{x}_t) \;=\; (\ell_{st} - \mu) + \hat{d}(\boldsymbol{x}_t) \;=\; (\ell_{st} - \mu) - v(\boldsymbol{x}_t) \;\;.$$

Here the inequality follows from the shortest path conditions on $\hat{d}()$. Therefore, by definition of ℓ_{st},

$$\begin{aligned} v(\boldsymbol{x}_t) \;&\leq\; v(\boldsymbol{x}_s) + \ell_{st} - \mu \\ &=\; v(\boldsymbol{x}_s) + \min_{s'} \{\boldsymbol{p}_{s'} \cdot (\boldsymbol{x}_t - \boldsymbol{x}_s) : \boldsymbol{x}_{s'} = \boldsymbol{x}_s\} - \mu \\ &\leq\; v(\boldsymbol{x}_s) + \boldsymbol{p}_s \cdot (\boldsymbol{x}_t - \boldsymbol{x}_s) - \mu \;\;. \end{aligned}$$

Hence, $v()$ is μ-approximate as desired.

Finally we require that $v()$ is minimum individually rational. So, take any other individually rational, μ-approximate virtual valuation $\omega()$. We must show that $v(\boldsymbol{x}_t) \leq \omega(\boldsymbol{x}_t)$ for every bundle $\boldsymbol{x}_t$. Now consider the shortest path tree T in H corresponding to $\hat{d}()$. If $(\boldsymbol{x}_t, \boldsymbol{z})$ is an arc in T (and at least one such arc exists) then $-\boldsymbol{p}_t \cdot \boldsymbol{x}_t = \hat{d}(\boldsymbol{x}_t)$. Thus

$$v(\boldsymbol{x}_t) - \boldsymbol{p}_t \cdot \boldsymbol{x}_t \;=\; (-\boldsymbol{p}_t \cdot \boldsymbol{x}_t) - \hat{d}(\boldsymbol{x}_t) \;=\; 0 \;\leq\; \omega(\boldsymbol{x}_t) - \boldsymbol{p}_t \cdot \boldsymbol{x}_t \;\;.$$

Here the inequality follows by the individual rationality of $\omega()$. Thus $v(\boldsymbol{x}_t) \leq \omega(\boldsymbol{x}_t)$. Now suppose that $v(\boldsymbol{x}_s) > \omega(\boldsymbol{x}_s)$ for some $\boldsymbol{x}_s$. We may take $\boldsymbol{x}_s$ to be the closest vertex to the root $\boldsymbol{z}$ in T with this property. We have seen that $\boldsymbol{x}_s$ cannot be a child of $\boldsymbol{z}$. So let $(\boldsymbol{x}_s, \boldsymbol{x}_t)$ be an arc in T. As $\boldsymbol{x}_t$ is closer to the root than $\boldsymbol{x}_s$, we know $v(\boldsymbol{x}_t) \leq \omega(\boldsymbol{x}_t)$. Then, as T is a shortest path tree, we have $\hat{d}(\boldsymbol{x}_s) = \hat{\ell}_{st} + \hat{d}(\boldsymbol{x}_t)$. Consequently $-v(\boldsymbol{x}_s) = \hat{\ell}_{st} - v(\boldsymbol{x}_t)$, and so

$$\omega(\boldsymbol{x}_t) \;\geq\; v(\boldsymbol{x}_t) \;=\; \hat{\ell}_{st} + v(\boldsymbol{x}_s) \;>\; \hat{\ell}_{st} + \omega(\boldsymbol{x}_s) \;\; .$$

But then

$$\omega(\boldsymbol{x}_t) \;>\; \omega(\boldsymbol{x}_s) + \ell_{st} - \mu \;=\; \omega(\boldsymbol{x}_s) + \min_{s'} \left\{ \boldsymbol{p}_{s'} \cdot (\boldsymbol{x}_t - \boldsymbol{x}_s) : \boldsymbol{x}_{s'} = \boldsymbol{x}_s \right\} - \mu \;\; .$$

It follows that there is at least one time period when $\boldsymbol{x}_s$ was selected in violation of the μ-optimality of $\omega()$. So $v()$ is a minimum individually rational, μ-approximate virtual valuation function. □

The minimum individually rational virtual valuation function allows us to obtain worst-case social welfare guarantees when revealed preference is used in mechanism design, see Section 2.5. For the best-case welfare guarantees, we are interested in finding the *maximum* virtual valuation function. In general, this need not exist as we may add an arbitrary constant to each bundle's valuation given by the minimum individually rational virtual valuation function. But, it does exist provided we have an upper bound on the valuation of at least one bundle. This is often the case. For example in a combinatorial auction if a bidder drops out of the auction at time $t+1$, then $\boldsymbol{p}_{t+1} \cdot \boldsymbol{x}_t$ is an upper bound on the value of bundle $\boldsymbol{x}_t$. Furthermore, in practice, bidders (and the auctioneer) often have (over)-estimates of the maximum possible value of some bundles.

So suppose we are given a set I and constraints of the form $v(\boldsymbol{x}_i) \leq \beta_i$ for each $i \in I$. Then there is a *unique* maximum μ-approximate virtual valuation function.

Theorem 2.4. *Given a set of constraints, the maximum μ-approximate virtual valuation function exists and can be found in polynomial time.*

Proof. Let $v(\boldsymbol{x}_i) \leq \beta_i$ for each $i \in I$. We construct a graph H from $\hat{G}$ by adding a source vertex $\boldsymbol{z}$ with arcs of length β_i from $\boldsymbol{z}$ to $\boldsymbol{x}_i$, for each $i \in I$. Since $\boldsymbol{z}$ has in-degree zero, H has no negative cycles because $\hat{G}$ does not. Denote by $\hat{d}()$ the shortest distance of every vertex *from* $\boldsymbol{z}$. We claim that setting $v(\boldsymbol{x}) = \hat{d}(\boldsymbol{x})$ gives us the desired maximum μ-approximate valuation function.

To prove this, we first begin by checking that it satisfies the upper-bound constraints. This is trivial, because for each $i \in I$ there is a path consisting of one arc of length β_i from $\boldsymbol{z}$ to $\boldsymbol{x}_i$. Thus the shortest path to $\boldsymbol{x}_i$ has length at most β_i. Second, the valuation function $v() = \hat{d}()$ is μ-approximate by the choice of arc length in $\hat{G}$. Third, we show that this valuation function is maximum. So, take any other μ-approximate virtual valuation $\omega()$ that satisfies the upper bound constraints I. We must show that $v(\boldsymbol{x}_t) \geq \omega(\boldsymbol{x}_t)$ for every bundle $\boldsymbol{x}_t$. For a contradiction, suppose that $P = \{\boldsymbol{z}, \boldsymbol{y}_1, \boldsymbol{y}_2, \ldots, \boldsymbol{y}_r\}$ is the shortest path from $\boldsymbol{z}$ to $\boldsymbol{y}_r$ in H and that $v(\boldsymbol{y}_r) < \omega(\boldsymbol{y}_r)$. Observe that the node adjacent to $\boldsymbol{z}$ on P must be $\boldsymbol{y}_1 = \boldsymbol{x}_i$ for some $i \in I$. Now because $\omega()$ is a μ-approximate valuation function, we have

$$\sum_{j=1}^{r-1} \omega(\boldsymbol{y}_{j+1}) \;\leq\; \sum_{j=1}^{r-1} \left(\omega(\boldsymbol{y}_j) + \ell_{\boldsymbol{y}_j,\boldsymbol{y}_{j+1}} - \mu\right) \;=\; \sum_{j=1}^{r-1} \left(\omega(\boldsymbol{y}_j) + \hat{\ell}_{\boldsymbol{y}_j,\boldsymbol{y}_{j+1}}\right) \;.$$

Cancelling terms produces

$$\omega(\boldsymbol{y}_r) \;\leq\; \omega(\boldsymbol{y}_1) + \sum_{j=1}^{r-1} \hat{\ell}_{\boldsymbol{y}_j,\boldsymbol{y}_{j+1}} \;\leq\; \beta_j + \sum_{j=1}^{r-1} \hat{\ell}_{\boldsymbol{y}_j,\boldsymbol{y}_{j+1}} \;=\; \hat{d}(\boldsymbol{y}_r) \;=\; v(\boldsymbol{y}_r) \;.$$

Here the second inequality follows by the facts that $\boldsymbol{y}_1 = \boldsymbol{x}_i$, for some $i \in I$, and $\omega()$ satisfies the upper bound constraint $\omega(\boldsymbol{x}_i) \leq \beta_i$. This contradicts the assumption that $v(\boldsymbol{y}_r) < \omega(\boldsymbol{y}_r)$. □

Notice that Theorem 2.4 does not guarantee that the maximum virtual valuation function is individually rational. For example, suppose $\beta_t = \boldsymbol{p}_t \cdot \boldsymbol{x}_t$, for all $1 \leq t \leq T$. Individual rationality then implies that $v(\boldsymbol{x}_t)$ must equal $\boldsymbol{p}_t \cdot \boldsymbol{x}_t$ for every bundle. In general, however, such a valuation function is not μ-approximate. In such cases no individually rational μ-approximate virtual valuation functions may exist that satisfy the upper bound constraints. On the other hand, suppose such a virtual valuation function does exist. Then the maximum μ-approximate virtual valuation function in Theorem 2.4 must be individually rational by maximality.

§2.5 Additive Relaxations to Revealed Preference Activity Rules

So far, we have focused upon how to test the degree of rationality reflected in a data set. Specifically, we saw in Theorem 2.2 that the minimum mean length of a cycle, $\mu(G)$, gives an exact and optimal goodness of fit measure for rationality. Furthermore, Theorem 2.3 explained how to quickly obtain the minimum individually rational valuation function that best fits the data.

Recall, however, that revealed preference is also used as a tool in mechanism design. In particular, we saw in Section 2.1 how revealed preference is used to impose bidding constraints in combinatorial auctions. We will now show how to apply the combinatorial arguments we have developed to create other relaxed revealed preference constraints.

Consider a combinatorial auction at time (round) t where our prior price-bundle bidding pairs are $\{(\boldsymbol{p}_1, \boldsymbol{x}_1), (\boldsymbol{p}_2, \boldsymbol{x}_2), \ldots, (\boldsymbol{p}_{t-1}, \boldsymbol{x}_{t-1})\}$. By Inequality (2.1.3) in section 2.1, rational bidding at time t implies that

$$v(\boldsymbol{x}_t) - \boldsymbol{p}_t \cdot \boldsymbol{x}_t \geq v(\boldsymbol{x}_s) - \boldsymbol{p}_t \cdot \boldsymbol{x}_s, \quad \text{for all } s < t.$$

Moreover, a necessary condition is then that $(\boldsymbol{p}_t - \boldsymbol{p}_s) \cdot \boldsymbol{x}_s \geq (\boldsymbol{p}_t - \boldsymbol{p}_s) \cdot \boldsymbol{x}_t$ and this can easily be checked by searching for negative length digons in the bidding graph induced by the first t bids. If such a cycle is found then the bid $(\boldsymbol{p}_t, \boldsymbol{x}_t)$ is not permitted by the auction mechanism.

The non-permittal of bids is clearly an extreme measure, and one that can lead to the exclusion of bidders from the auction even when they still have bids they wish to make. In this respect, it may be desirable for the mechanism to use a relaxed set of revealed preference bidding rules. The natural approach is to insist not upon strictly rational bidders but rather just upon approximately rational bidders. Specifically, the auction mechanism may (dynamically) select a desired degree ϵ of rationality. This requires that at time t,

$$v(\boldsymbol{x}_t) - \boldsymbol{p}_t \cdot \boldsymbol{x}_t \geq v(\boldsymbol{x}_s) - \boldsymbol{p}_t \cdot \boldsymbol{x}_s - \epsilon, \quad \text{for all } s < t.$$

A necessary condition then is $(\boldsymbol{p}_t - \boldsymbol{p}_s) \cdot \boldsymbol{x}_s \geq (\boldsymbol{p}_t - \boldsymbol{p}_s) \cdot \boldsymbol{x}_t - 2\epsilon$, and we can test this *relaxed* WARP-based bidding rule by insisting that every digon has mean length at least $-\epsilon$.

Similarly, the *relaxed* KARP-based bidding rule is

$$(\boldsymbol{p}_k - \boldsymbol{p}_0) \cdot \boldsymbol{x}_0 \ \geq \ \sum_{i=1}^{k} (\boldsymbol{p}_i - \boldsymbol{p}_{i-1}) \cdot \boldsymbol{x}_i - (k+1) \cdot \epsilon \tag{2.5.1}$$

The *relaxed* GARP-based bidding rule applies the relaxed KARP-based bidding rule for every choice of k. The imposition of the relaxed GARP-based bidding rule ensures approximate rationality.

Theorem 2.5. *A set of price-bid pairings* $\{(\boldsymbol{p}_t, \boldsymbol{x}_t) : 1 \leq t \leq T\}$ *has a corresponding* ϵ*-approximate individually rational virtual valuation function if and only if it satisfies the relaxed* GARP*-based bidding rule.*

Proof. Suppose the relaxed GARP-based bidding rule is satisfied. By Theorem 2.2, it suffices to show that the minimum mean cycle in the bidding graph with arc lengths ℓ is at least $-\epsilon$. So take any collection $\{\boldsymbol{x}_i\}_{i=1}^k$ of bundles. Let t_i be the time when $\ell_{\boldsymbol{x}_i, \boldsymbol{x}_{i+1}}$ was minimized, and let $\boldsymbol{p}_i := \boldsymbol{p}_{t_i}$. Then we have

$$\begin{aligned} -(k+1) \cdot \epsilon \ &\leq \ (\boldsymbol{p}_k - \boldsymbol{p}_0) \cdot \boldsymbol{x}_0 - \sum\nolimits_{i=1}^{k} (\boldsymbol{p}_i - \boldsymbol{p}_{i-1}) \cdot \boldsymbol{x}_i \\ &= \sum\nolimits_{i=0}^{k} \boldsymbol{p}_i \cdot (\boldsymbol{x}_{i+1} - \boldsymbol{x}_i) \\ &= \sum\nolimits_{i=0}^{k} \ell_{\boldsymbol{x}_i, \boldsymbol{x}_{i+1}} \end{aligned}$$

Here, the inequality follows because the relaxed GARP-based bidding rule is satisfied. (Again the subscripts are taken modulo $k+1$.) Since, the corresponding cycle contains $k+1$ arcs, we see that the length of the minimum mean cycle is at least $-\epsilon$.

Conversely, if the bidding data has a corresponding ϵ-approximate individually rational virtual valuation function then the relaxed bidding rules are satisfied. □

Theorem 2.5 tells us that imposing the relaxed GARP-based bidding rule ensures approximate rationality. But, in practice, even WARP-based bidding rules are often confusing to real bidders. There is likely therefore to be some resistance to the idea of imposing the whole gamut of GARP-based bidding rules. We believe that this combinatorial view of revealed preference, where the bidding rules can be tested via cycle examination, will eradicate some of the confusion. However, for simplicity, there is some worth in quantitatively examining the consequences of imposing a weaker relaxed KARP-based bidding rule rather than the GARP-based bidding rule. To test for the relaxed KARP-based bidding rules, we simply have to examine cycles of length at most $k+1$. Now suppose the KARP-based bidding rules are satisfied. By finding the $\mu(G)$ in the bidding graph we can still obtain the best-fit additive approximation guarantee, but we no longer have that this guarantee is ϵ. We can still, though, prove a strong additive approximation guarantee even for small values of k. To do this we need the following result.

Theorem 2.6. *Given a complete directed graph G with arc lengths ℓ. If every cycle of cardinality at most $k+1$ has non-negative length then the minimum mean length of a cycle is at least $-\frac{\ell^{\max}}{k}$, where $\ell^{\max} = \max_{e \in E(G)} |\ell_e|$.*

Proof. Take any cycle C with cardinality $|C| > k+1$. Let the arcs of C be $\{e_1, e_2, \ldots, e_{|C|}\}$ in order. Then

$$\sum_{i=1}^{|C|} \sum_{j=i}^{i+k-1} \ell_{e_j} \;=\; k \cdot \sum_{i=1}^{|C|} \ell_{e_i} \;=\; k \cdot \ell(C) \;=\; k \cdot |C| \cdot \frac{\ell(C)}{|C|} \;. \tag{2.5.2}$$

Above, the inner summation is taken modulo $|C|$. On the other hand take any path segment $P = \{e_i, e_{i+1}, \ldots, e_{i+k-1}\}$, where again the subscript summation is modulo $|C|$. Because the graph is complete and the maximum arc length is $\ell^{\max}$, the length of P is at least $-\ell^{\max}$. Otherwise, we have a negative length cycle of cardinality $k+1$ by adding to P the arc from the head vertex of e_{i+k-1} to the tail vertex of e_i. Thus,

$$\sum_{i=1}^{|C|} \sum_{j=i}^{i+k-1} \ell_{e_j} \geq -|C| \cdot \ell^{\max} \ . \tag{2.5.3}$$

Combining Equalities (2.5.2) and Inequality (2.5.3) gives that $\frac{\ell(C)}{|C|} \geq -\frac{\ell^{\max}}{k}$. As every cycle of cardinality at most $k+1$ has non-negative mean length, this implies that the minimum mean length of any cycle in G is at least $-\frac{\ell^{\max}}{k}$. □

This result is important as it allows us to bound the degree of rationality that must arise whenever we impose the relaxed KARP-based bidding rule.

Corollary 2.1. *Given a set of price-bid pairings $\{(\boldsymbol{p}_t, \boldsymbol{x}_t) : 1 \leq t \leq T\}$ that satisfy the relaxed* KARP*-based bidding rule, there is a $(\frac{b^{\max}}{k} + \epsilon)$-approximate individually rational virtual valuation function, where $b^{\max}$ is the maximum bid made by the bidder during the auction.*

Proof. The relaxed KARP-based bidding rule (2.5.1) implies that every cycle of cardinality at most $k+1$ in the bidding graph G has mean length at least $-\epsilon$. Let G' be the modified graph with arc lengths $\ell'_{\boldsymbol{x}_s,\boldsymbol{x}_t} := \ell_{\boldsymbol{x}_s,\boldsymbol{x}_t} + \epsilon$. Then every cycle in G' of cardinality at most $k+1$ has non-negative length. By Theorem 2.6, the minimum mean length of a cycle in G' is then at most $\frac{(\ell')^{\max}}{k}$. Furthermore, $(\ell')^{max} = \ell^{max} + \epsilon \leq b^{\max} + \epsilon$. Theorems 2.2 and 2.3 then guarantee the existence of a $(\frac{b^{\max}}{k} + \epsilon)$-approximate individually rational virtual valuation function. $\square$

One may ask whether the additive approximation guarantee in Corollary 2.1 can be improved. The answer is *no*; Theorem 2.6 is tight.

Lemma 2.1. *There is a graph G where each cycle of cardinality at most $k+1$ has non-negative length and the minimum mean length of a cycle is $-\ell^{\max}/k$.*

Proof. Let G be a complete directed graph with vertex set $V = \{v_1, v_2, \ldots, v_n\}$. We will define arc lengths ℓ such that all $(k+1)$-cycles in G have non-negative length, but the minimum mean length of a cycle is $-\frac{\ell^{\max}}{k}$. First consider the cycle $C_0 = \{v_1, v_2, \ldots, v_{k+2}, v_1\}$. Give each arc in C_0 a length $-\frac{\ell^{\max}}{k}$. Thus C_0 has cardinality $k+2$ and mean length $-\frac{\ell^{\max}}{k}$. Now let every other arc e have length $\ell^{\max}$. It immediately follows that the only cycle in G with negative length is C_0. Thus, all cycles of length at most $k+1$ have non-negative length, but the minimum mean length of a cycle is $-\frac{\ell^{\max}}{k}$, as desired. $\square$

§2.6 Alternate Bidding Rules

Interestingly other bidding rules used in practice or proposed in the literature can be viewed in the graphical framework. For example, bid withdrawals correspond to vertex deletion in the bidding graph, whilst budget constraints and the Afriat Efficiency Index can be formulated in terms of arc-deletion. We briefly describe these applications here. See Chapter 3 for a more detailed study of the computational complexity of bid-withdrawal rules.

Revealed Preference with Budgets. Recall that, in Section 2.1, we have assumed that, in the quasilinear model, bidders have no budgetary constraints. This is not a natural assumption. Harsha et al. [24] explain how to implement budgeted revealed preference in a combinatorial auction. Their method applies to the case when the fixed budget B is unknown to the auction mechanism. To do this, upper and lower bounds on feasible budgets are maintained dynamically via a linear program. It is also straightforward to do this combinatorially using edge-deletion in the bidding graph; we omit the details as the process resembles that of the following subsection.

The Afriat Efficiency Index. Recall that to determine the Afriat Efficiency Index we reveal $\boldsymbol{x}_t \succeq \boldsymbol{y}$ only if $\boldsymbol{p}_t \cdot \boldsymbol{y} \leq \lambda \cdot \boldsymbol{p}_t \cdot \boldsymbol{x}_t$ where $\lambda < 1$. This is equivalent, in Afriat's original setting, to removing from the graph any arc $(\boldsymbol{x}_t, \boldsymbol{x}_s)$ for which $\boldsymbol{p}_t \cdot \boldsymbol{x}_s > \lambda \cdot \boldsymbol{p}_t \cdot \boldsymbol{x}_t$. Of course, for the application of combinatorial auctions, we assume quasi-linear utilities. Therefore, the appropriate implementation is to remove any arc $(\boldsymbol{x}_t, \boldsymbol{x}_s)$ for which

$$v(\boldsymbol{x}_s) - \boldsymbol{p}_t \boldsymbol{x}_s > \lambda \cdot (v(\boldsymbol{x}_t) - \boldsymbol{p}_t \boldsymbol{x}_t) \ .$$

How, though, can we implement this rule as $v()$ is unknown? We can simply apply the techniques of Section 2.3 and use for v the minimum individually rational virtual valuation function. We can now determine the best choice of λ that gives a predetermined, ϵ additive approximation guarantee ϵ. This can easily be computed exactly by bisection search over the set of arcs, as each arc a has its own critical value λ_a at which it will be removed. The optimal choice arises at the point where the minimum mean cycle in the bidding graph rises above $-\epsilon$. When $\epsilon = 0$, the corresponding choice of λ is the anolog of the Afriat Efficiency Index.

Revealed Preference with Bid Withdrawals. Some iterative multi-item auctions allow for bid withdrawals, most notably the simultaneous multi-round auction (SMRA). Bid withdrawals may easily be implemented along with revealed preference bidding rules. At time t, a bid withdrawal corresponds to the removal of (a copy of) a vertex $\boldsymbol{x}_s$, where $s < t$. This may be important strategically. To see this, suppose the bid $\boldsymbol{x}_t$ is invalid under the KARP-based bidding rules because it would induce a negative cycle of cardinality at most $k+1$ in the bidding graph on $\{\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_t\}$. If $\boldsymbol{x}_s$ lies on all such negative cycles then $\boldsymbol{x}_t$ becomes a valid bid after the withdrawal of $\boldsymbol{x}_s$. Because auctions typically restrict the total number of bid withdrawals allowed, the optimal application of bid withdrawals correspond to the problem of finding small hitting sets for the negative length cycles of cardinality at most $k+1$.

Chapter 3

A Combinatorial Measure of Rationality in the General Setting: Perfect Graphs and Intractability Results

In Section 2.6, we introduced the problem of finding small hitting sets for cycles in the preference data. In this chapter, we consider the problem in Afriat's original setting of an agent in a commodity market, and study its computational complexity. We show that, unless the market has only two commodities, the problem is NP-hard. Therefore, such a combinatorial measure of consistency may be difficult to implement in some cases.

Recall that we have observed a collection of *consumer data* $\{(\boldsymbol{p}_1, \boldsymbol{x}_1), (\boldsymbol{p}_2, \boldsymbol{x}_2), \ldots, (\boldsymbol{p}_m, \boldsymbol{x}_m)\}$, where each pair $(\boldsymbol{p}_i, \boldsymbol{x}_i)$ denotes the fact that the consumer purchased the bundle of goods $\boldsymbol{x}_i \in \mathbf{R}^n$ when the prices were $\boldsymbol{p}_i \in \mathbf{R}^n$. For clarity of presentation, we will assume that all the chosen bundles are distinct and that all revealed preferences are strict (no ties). We can represent the preferences revealed by the consumer data via a directed graph, $D_{\succeq} = (V, A)$. This directed *revealed preference graph* contains a vertex $\boldsymbol{x}_i \in V$ for each data-pair $(\boldsymbol{p}_i, \boldsymbol{x}_i)$, and an arc from $\boldsymbol{x}_i$ to $\boldsymbol{x}_j$ if and only if $\boldsymbol{x}_i \succeq \boldsymbol{x}_j$. Observe that GARP holds *if and only if* the revealed preference graph is acyclic. Consequently, Afriat's theorem implies that the consumer is rational if and only if $D_{\succeq}$ contains no directed cycles. Observe also that this graph is similar, but different, to the the *bidding graph* constructed in Chapter 2.

For example, Figure 3.0.1 displays visually two sets of consumer data. Each bundle $\boldsymbol{x}_i$ is paired with its price vector $\boldsymbol{p}_i$, and a dotted line is drawn through $\boldsymbol{x}_i$ perpendicular to $\boldsymbol{p}_i$. Note that $\boldsymbol{p}_i\boldsymbol{x}_i \geq \boldsymbol{p}_i\boldsymbol{y}$ if and only if $\boldsymbol{y}$ lies on the opposite side of the dotted line to the drawing of $\boldsymbol{p}_i$. Hence, for the first consumer (left), we have $\boldsymbol{x}_3 \succeq \boldsymbol{x}_2$, $\boldsymbol{x}_3 \succeq \boldsymbol{x}_1$ and $\boldsymbol{x}_2 \succeq \boldsymbol{x}_1$. This produces an acyclic revealed preference graph $D_{\succeq}$ and, therefore, her behaviour can be rationalized. On the otherhand, the second consumer (right) reveals $\boldsymbol{x}_3 \succeq \boldsymbol{x}_2 \succeq \boldsymbol{x}_3$. This produces a directed 2-cycle in $D_{\succeq}$ and, so, her behaviour cannot be rationalised.

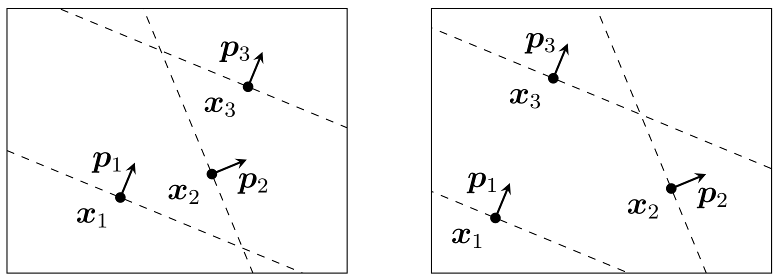

Figure 3.0.1: A rational consumer and an irrational consumer.

We have seen that graph acyclicity can be used to provide a test for consumer rationality. However such a test is binary and, in practice, leads to the immediate conclusion of irrationality, as observed data sets typically induce cycles in the revealed preference graph. Consequently, there has been a large body of experimental and theoretical work designed to measure how close to rational the behaviour of a consumer is. Examples include measurements based upon best-fit perturbation errors (*e.g.* Afriat [2] and Varian [40]), measurements based upon counting the number of rationality violations present in the data (*e.g.* Swofford and Whitney [38] and Famulari [17]), and measurements based upon the maximum size of a *rational* subset of the data (*e.g.* Houtman and Maks [27]). Gross [19] provides a review and analysis of some of these measures. Recently new measures have been designed by Echenique et al. [16], Apesteguia and Ballester [4], and Dean and Martin [11].

Combinatorially, perhaps the most natural measure is simply to count the number of "irrational" purchases. That is, what is the minimum number of data-points whose removal induces a rational set of data? The associated decision problem is called the CONSUMER RATIONALITY problem.

CONSUMER RATIONALITY

Instance: Consumer data $(\boldsymbol{p}_1, \boldsymbol{x}_1), \ldots, (\boldsymbol{p}_m, \boldsymbol{x}_m) \in \mathbf{R}^n \times \mathbf{R}^n$, and an integer k.

Problem: Is there a sub-collection of at most k data points whose removal produces a data set satisfying GARP?

We note that this CONSUMER RATIONALITY problem is dual to the measure of Houtman and Maks [27]. Using the graphical representation, it can be seen that the consumer rationality problem is a special case of the DIRECTED FEEDBACK VERTEX SET problem. In fact, as we explain in Section 3.1, when there are many goods, the two problems are equivalent. However, the consumer rationality problem becomes easier to approximate as the number of commodities falls. Indeed, the main contribution of this paper is to obtain an exact threshold on the number of commodities that separates easy cases (polynomial) and hard cases (NP-complete). In particular, we prove the problem is polytime solvable for a two-commodity market (Section 3.3), but that it is NP-complete for a three-commodity market (Section 3.4).

§3.1 The General Case: Many Commodities

In this section we show that the CONSUMER RATIONALITY problem in full generality is computationally equivalent to the DIRECTED FEEDBACK VERTEX SET (DFVS) problem.

DIRECTED FEEDBACK VERTEX SET

Instance: A directed graph $D = (V, A)$, and an integer k.

Problem: Is there a set S of at most k vertices such that the induced subgraph $D[V \setminus S]$ is acyclic? (Such a set S is called a *feedback vertex set.*)

First, observe that the CONSUMER RATIONALITY PROBLEM is a special case of the DIRECTED FEEDBACK VERTEX SET PROBLEM: we have seen that the dataset is rationalizable if and only if the preference graph is acyclic. Thus, the minimum feedback vertex set in the preference graph $D_{\succeq}$ clearly corresponds to the minimum number of data points that must be removed to create a rationalizable data-set.

On the other hand, provided the number of commodities is large, DFVS is a special case of the CONSUMER RATIONALITY PROBLEM. Specifically, Deb and Pai [12] show that for any directed graph D there is a data-set on $m = n$ commodities whose preference graph is $D_{\succeq} = D$; for completeness, we include the short proof of this result.

Lemma 3.1. *[12] Given sufficiently many commodities, we may construct any digraph as a preference graph.*

Proof. Let D be any digraph on n nodes. We will construct n pairs in $\mathbf{R}^n \times \mathbf{R}^n$ such that $D_{\succeq} \cong D$. Denote $\boldsymbol{p}^i = (p^i_1, \ldots, p^i_n)$, and set $p^i_i = 1$, $p^i_j = 0$ for $j \neq i$. Similarly, denote $\boldsymbol{x}^i = (x^i_1, \ldots, x^i_n)$, and set $x^i_j = 1$ if $i = j$, 0 if $(i, j) \in D$, and 2 if $(i, j) \notin D$. We then have, $\boldsymbol{p}_i \cdot \boldsymbol{x}_i = 1$, $\boldsymbol{p}_i \cdot \boldsymbol{x}_j = 0$ if we want an arc from i to j, and $\boldsymbol{p}_i \cdot \boldsymbol{x}_j = 2$ if we do not want an arc, as desired. $\square$

It follows that any lower and upper bounds on approximation for (the optimization version of) DFVS immediately apply to (the optimization version of) the CONSUMER RATIONALITY problem. The exact hardness of approximation for DFVS is not known. The best upper bound is due to Seymour [36] who gave an $O(\log n \log\log n)$ approximation algorithm. With respect to lower bounds, the DIRECTED FEEDBACK VERTEX SET problem is NP-complete [28]. Furthermore, as we will see in Section 3.2, the CONSUMER RATIONALITY problem is at least as hard to approximate as VERTEX COVER. It follows that DFVS problem cannot be approximated to within a factor 1.36 [13] unless $P = NP$. Also, assuming the Unique Games Conjecture [30], the minimum directed feedback vertex set cannot be approximated to within any constant factor [23, 37].

Lemma 3.1 shows the equivalence with DIRECTED FEEDBACK VERTEX SET applies when the number of commodities is at least the size of the data-set. However, Deb and Pai [12] also show that for an m-commodity market, there exists a directed graph on $O(2^m)$ vertices that cannot be realised as a preference graph. This suggests that the hardness of the consumer rationality problem may vary with the quantity of goods. Indeed, we now prove that this is the case.

§3.2 Two-Commodity Markets and the Vertex Cover Problem

We begin by outlining the basic approach to proving polynomial solvability for two goods. As described, the CONSUMER RATIONALITY problem is a special case of DVFS. For two goods, however, rather than considering all directed cycles, it is sufficient to find a vertex hitting set for the set of *digons* (directed cycles consisting of two arcs). The resulting problem can be solved by finding a minimum vertex cover in a corresponding auxiliary undirected graph. The VERTEX COVER problem is, of course, itself hard [13]. But we prove that the auxiliary undirected graph is perfect, and VERTEX COVER is polytime solvable in perfect graphs.

So, our first step is to show that it suffices to hit only digons. Specifically, we prove that every vertex-minimal cycle in the revealed preference graph $D_{\succeq}$ is a digon. This fact corresponds to the result that for two goods the *Weak Axiom of Revealed Preference* is equivalent to the *Generalised Axiom of Revealed Preference.* This equivalence was noted by Samuelson [35] and formally proven by Rose [33] in 1958; for a recent structurally motivated proof see [25]. For completeness, and to illustrate some of the notation and techniques required in this paper, we present a short geometric proof here.

We begin with the required notation. Let $\boldsymbol{x} = (x_1, x_2) \in \mathbf{R}^2$, and define

$$\boldsymbol{x}^{\searrow} := \{(y_1, y_2) \in \mathbf{R}^2 : y_1 \geq x_1,\ y_2 \leq x_2\} \ ,$$

i.e. the points which lie "below and to the right" of $\boldsymbol{x}$ in the plane. Define $\boldsymbol{x}^{\nwarrow}$, $\boldsymbol{x}^{\nearrow}$ and $\boldsymbol{x}^{\swarrow}$ similarly. In addition, define $\boldsymbol{x}^{\searrow\!\!\searrow}$ $\boldsymbol{x}^{\nwarrow\!\!\nwarrow}$, $\boldsymbol{x}^{\nearrow\!\!\nearrow}$ and $\boldsymbol{x}^{\swarrow\!\!\swarrow}$ by replacing the inequalities with strict inequalities. Furthermore, if ℓ is a line in the plane of non-positive slope which intersects the positive quadrant, we say a point *lies below* ℓ if it lies in the same closed half-plane as the origin. For each data pair $(\boldsymbol{p}_i, \boldsymbol{x}_i)$, we define ℓ_i to be the line through $\boldsymbol{x}_i$ perpendicular to $\boldsymbol{p}_i$. Hence, in our setting $\boldsymbol{x}_i \succeq \boldsymbol{x}_j$ if and only if $\boldsymbol{x}_j$ lies below ℓ_i. Note that, if $\boldsymbol{x}_i \succeq \boldsymbol{x}_j$, then we may not have $\boldsymbol{x}_j \in \boldsymbol{x}_i^{\nearrow\!\!\nearrow}$ since $\boldsymbol{p}_i$ is non-negative.

Lemma 3.2. *[33] For two commodities, every minimal cycle is a digon.*

Proof. Let $C_k = \{\boldsymbol{x}_0, \boldsymbol{x}_1, \ldots, \boldsymbol{x}_{k-1}\}$, listed in order, be a vertex-minimal directed cycle in $D_{\succeq}$. Suppose, for a contradiction, that $k \geq 3$. By minimality, the cycle C_k is chordless, therefore, $\boldsymbol{x}_i \succeq \boldsymbol{x}_j$ if and only if $j = i + 1 \pmod{k}$. (Henceforth, we will often assume without statement that indices are taken modulo k. Furthermore, "left" will stand for the negative x direction.) Without loss of generality, suppose $\boldsymbol{x}_i$ is the leftmost bundle – or one of them. Since $\boldsymbol{x}_i \succeq \boldsymbol{x}_{i+1}$, we have that $\boldsymbol{x}_{i+1}$ must fall in $\boldsymbol{x}_i^{\searrow}$. We claim that ℓ_i must be steeper than ℓ_{i+1}. To see this, suppose this is not true. Then, as shown in Figure 3.2.1(a), ℓ_{i+1} must intersect the line ℓ_i strictly to the left of $\boldsymbol{x}_i$. If not, $\boldsymbol{x}_{i+1} \succeq \boldsymbol{x}_i$. Now $\boldsymbol{x}_{i+2}$ lies under ℓ_{i+1} but not under ℓ_i, but this implies that $\boldsymbol{x}_{i+2}$ lies strictly to the left of $\boldsymbol{x}_i$ as illustrated. This gives the desired contradiction. Hence, ℓ_i must be steeper than ℓ_{i+1}. This situation is

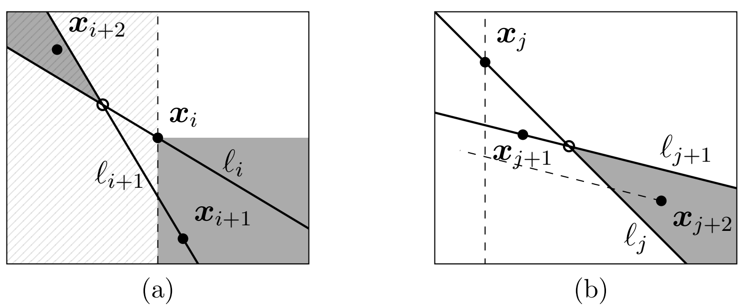

Figure 3.2.1: Leftmost 2-commodity bundles on a cycle.

illustrated in Figure 3.2.1(b) where we set $j = i$. We claim the following:

Claim 3.1. *Suppose $\boldsymbol{x}_{j+1} \in \boldsymbol{x}_j^{\searrow}$ and ℓ_j is steeper than ℓ_{j+1}, then we must have that $\boldsymbol{x}_{j+2} \in \boldsymbol{x}_{j+1}^{\searrow}$ and that ℓ_{j+1} is steeper than ℓ_{j+2}.*

As shown in Figure 3.2.1(b), because ℓ_j is steeper than ℓ_{j+1}, we must have $\boldsymbol{x}_{j+2} \in \boldsymbol{x}_{j+1}^{\searrow}$. It remains to show that ℓ_{j+1} is steeper than ℓ_{j+2}. Suppose not, then, since $\boldsymbol{x}_{j+1}$ must fall above ℓ_{j+2}, the (highlighted) point where ℓ_j and ℓ_{j+1} meet must also fall above ℓ_{j+2}. Thus, the region which falls above both ℓ_j and ℓ_{j+1} cannot intersect the region below ℓ_{j+2}. Therefore, there is no valid position for $\boldsymbol{x}_{j+3}$. Consequently, ℓ_{j+1} must be steeper than ℓ_{j+2}, as desired.

Hence, by induction, for every $0 \leq j \leq k-1$, we have that ℓ_j is steeper than ℓ_{j+1} and that $\boldsymbol{x}_{j+1} \in \boldsymbol{x}_j^{\searrow}$, where our base case is $j = i$. However, this cannot hold for $j = i-1$, since $\boldsymbol{x}_i$ is the leftmost point in the cycle, amounting to a contradiction, and refuting the assumption that there existed a minimal cycle on at least 3 vertices. □

Lemma 3.2 implies that a vertex set that intersects every digon will also intersect each directed cycle of any length. Hence, to solve the CONSUMER RATIONALITY problem for two goods, it suffices to find a minimum cardinality hitting vertex set for the digons of $D_{\succeq}$. We can do this by transforming the problem into one of finding a minimum vertex cover in an undirected graph. Recall the VERTEX COVER problem is:

VERTEX COVER

INSTANCE: Given an undirected graph $G = (V, E)$ and an integer k.

PROBLEM: Is there a set S of at most k vertices such that every edge has an endpoint in S?

The transformation is then as follows: given the directed revealed preference graph $D_{\succeq}$ we create an *auxiliary undirected graph* $G_{\succeq}$. The vertex set $V(G_{\succeq}) = V(D_{\succeq})$ so the undirected graph also has a vertex for each bundle $\boldsymbol{x}_i$. There is an edge $(\boldsymbol{x}_i, \boldsymbol{x}_j)$ in $G_{\succeq}$ if and only if $\boldsymbol{x}_i$ and $\boldsymbol{x}_j$ induce a digon in $D_{\succeq}$. It is easy to verify that a vertex cover in $G_{\succeq}$ corresponds to a hitting set for digons of $D_{\succeq}$.

Let's see some simple examples for the auxiliary graph $G_{\succeq}$. First consider Figure 3.2.2(a), where bundles are placed on a concave curve. Now every pair of vertices $\boldsymbol{x}_i$ and $\boldsymbol{x}_j$ induce a digon in $D_{\succeq}$. Thus $G_{\succeq}$ is an undirected clique. Now consider Figure 3.2.2(b). The vertices on the left induce a directed path in $D_{\succeq}$; the vertices along the bottom also induce a directed path in $D_{\succeq}$. However each pair consisting of one vertex on the left and one vertex on the bottom induce a digon in $D_{\succeq}$. Thus $G_{\succeq}$ is a complete bipartite graph.

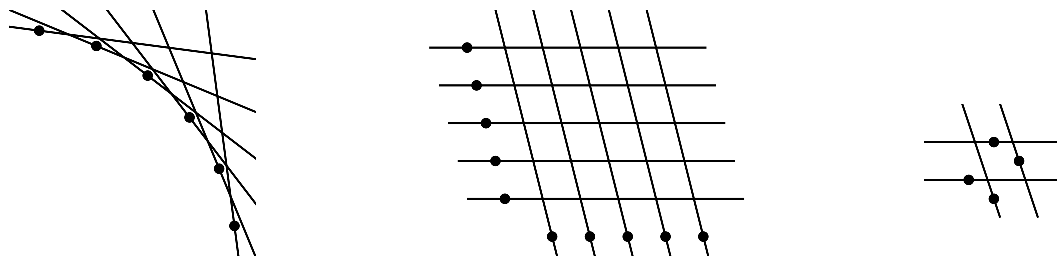

(a) A Complete Graph (b) A Complete Bipartite Graph (c) $K_2 \cup K_2$

Figure 3.2.2: Examples of the Auxiliary Undirected Graph.

§3.3 Auxiliary Graphs in 2-Commodity Market are Perfect Graphs

An undirected graph G is perfect if the chromatic number of any induced subgraph is equal to the cardinality of the maximum clique in the subgraph. In 1961, Berge [7] made the famous conjecture that an undirected graph is perfect if and only if it contains neither an odd length hole nor an odd length antihole. Here a *hole* is a chordless cycle with at least four vertices. An *antihole* is the complement of a chordless cycle with at least four vertices. Berge's conjecture was finally proven by Chudnovsky, Robertson, Seymour and Thomas [9] in 2006.

Theorem 3.1 (The Strong Perfect Graph Theorem [9]). *An undirected graph is perfect if and only if it contains no odd holes and no odd antiholes.*

There are many important classes of perfect graphs, for example, cliques, bipartitie graphs, chordal graphs, line graphs of bipartite graphs, and comparability graphs.[1] Interestingly, we now show that the class of 2D auxiliary revealed preference graphs are also perfect. To prove this, we will need the following geometric lemma, but first, we introduce the required notation.

Lemma 3.3. *Let $\{x_i, x_j, x_k\}$, listed in order, be an induced path in the 2D auxiliary revealed preference graph $G_{\succeq}$. If $\boldsymbol{x}_i \in \boldsymbol{x}_j^{\nwarrow}$ then $\boldsymbol{x}_k \in \boldsymbol{x}_j^{\nwarrow}$. (Similarly, if $\boldsymbol{x}_i \in \boldsymbol{x}_j^{\searrow}$ then $\boldsymbol{x}_k \in \boldsymbol{x}_j^{\searrow}$.)*

Proof. Recall the assumption that the bundles distinct, that is, $\boldsymbol{x}_i \neq \boldsymbol{x}_j$ for all $i \neq j$. Because $\{x_i, x_j\}$ is an edge in the auxiliary undirected graph $G_{\succeq}$, we know that $\boldsymbol{x}_i \succeq \boldsymbol{x}_j$ and $\boldsymbol{x}_j \succeq \boldsymbol{x}_i$. Therefore it cannot be the case that $\boldsymbol{x}_i \in \boldsymbol{x}_j^{\nearrow\nearrow}$ or $\boldsymbol{x}_j \in \boldsymbol{x}_i^{\nearrow\nearrow}$. Thus, either $\boldsymbol{x}_j \in \boldsymbol{x}_i^{\nwarrow}$ or $\boldsymbol{x}_j \in \boldsymbol{x}_i^{\searrow}$, but not both. Similarly, because $\{x_j, x_k\}$ is an edge in $G_{\succeq}$, either $\boldsymbol{x}_k \in \boldsymbol{x}_j^{\nwarrow}$ or $\boldsymbol{x}_k \in \boldsymbol{x}_j^{\searrow}$.

Now, without loss of generality, let $\boldsymbol{x}_i \in \boldsymbol{x}_j^{\nwarrow}$. For a contradiction, assume that $\boldsymbol{x}_k \in \boldsymbol{x}_j^{\searrow}$. Hence, we have $\boldsymbol{x}_j \in \boldsymbol{x}_i^{\searrow} \cap \boldsymbol{x}_k^{\nwarrow}$. Suppose $\boldsymbol{x}_j$ lies strictly below the line $\ell_{i,k}$ through $\boldsymbol{x}_i$ and $\boldsymbol{x}_k$. But then we cannot have both $\boldsymbol{x}_j \succeq \boldsymbol{x}_i$ and $\boldsymbol{x}_j \succeq \boldsymbol{x}_k$. This is because the line ℓ_j must cross the segment of $\ell_{i,k}$ between $\boldsymbol{x}_i$ and $\boldsymbol{x}_k$ if it is to induce either of the two preferences. Thus, the line ℓ_j separates $\boldsymbol{x}_i$ and $\boldsymbol{x}_k$ and, so, at most one of bundles can lie below the line. This is illustrated in Figure 3.3.1(a).

On the other hand, suppose $\boldsymbol{x}_j$ lies on or above the line $\ell_{i,k}$ through $\boldsymbol{x}_i$ and $\boldsymbol{x}_k$. Now we know that $\boldsymbol{x}_i \succeq \boldsymbol{x}_j$. This implies that $\boldsymbol{x}_i \succeq \boldsymbol{x}_k$, as illustrated in Figure 3.3.1(b). Furthermore, we know that $\boldsymbol{x}_k \succeq \boldsymbol{x}_j$ which implies that $\boldsymbol{x}_k \succeq \boldsymbol{x}_i$. Thus $\{x_i, x_k\}$ is an edge in $G_{\succeq}$. This

[1] By the (Weak) Perfect Graph Theorem [31], the complements of these classes of graphs are also perfect.

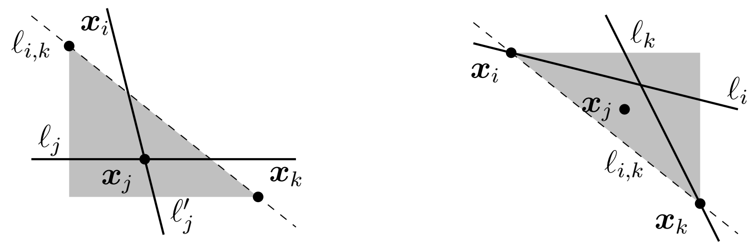

(a) $\boldsymbol{x}_j$ below $\ell_{i,k}$ (b) $\boldsymbol{x}_j$ above or on $\ell_{i,k}$

Figure 3.3.1: Induced path on three vertices.

contradicts the fact that $\{x_i, x_j, x_k\}$ is an induced path. □

Lemma 3.4. *The 2D auxiliary revealed preference graph $G_{\succeq}$ contains no odd holes on at least 5 vertices.*

Proof. Take a hole $C_k = \{\boldsymbol{x}_0, \boldsymbol{x}_1, \ldots, \boldsymbol{x}_{k-1}\}$, listed in order, where $k \geq 5$ is odd. For any $0 \leq i \leq k-1$, the three vertices $\{\boldsymbol{x}_{i-1}, \boldsymbol{x}_i, \boldsymbol{x}_{i+1}\}$ induce a path in $G_{\succeq}$. Consequently, by Lemma 3.3, either both $\boldsymbol{x}_{i-1}$ and $\boldsymbol{x}_{i+1}$ are in $\boldsymbol{x}_i^{\nwarrow}$ or both $\boldsymbol{x}_{i-1}$ and $\boldsymbol{x}_{i+1}$ are in $\boldsymbol{x}_i^{\searrow}$. In the former case, colour $\boldsymbol{x}_i$ yellow. In the latter case, colour $\boldsymbol{x}_i$ red. Thus we obtain a 2-coloring of C_k. Since k is odd, there must be two adjacent vertices, $\boldsymbol{x}_i$ and $\boldsymbol{x}_{i+1}$, with the same colour. Without loss of generality, let both vertices be yellow. Thus, $\boldsymbol{x}_{i+1}$ is $\boldsymbol{x}_i^{\nwarrow}$ and $\boldsymbol{x}_i$ is in $\boldsymbol{x}_{i+1}^{\nwarrow}$. This contradicts the distinctness of $\boldsymbol{x}_i$ and $\boldsymbol{x}_{i+1}$. □

We remark that the parity condition in Lemma 3.4 is necessary. To see this consider the example in Figure 3.3.2 which produces an even hole on six vertices. Specifically, the only mutually adjacent pairs are the $(\boldsymbol{x}_i, \boldsymbol{x}_{i+1})$ pairs, with indices taken modulo 6.

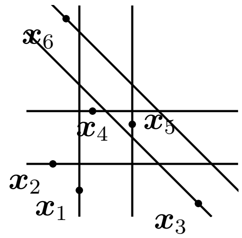

Figure 3.3.2: Construction for C_6

Lemma 3.5. *The 2D auxiliary revealed preference graph $G_{\succeq}$ contains no antiholes on at least 5 vertices.*

Proof. Note that the complement of an odd hole on five vertices is also an odd hole. Thus, by Lemma 3.4, the graph $G_{\succeq}$ may not contain an antihole on five vertices.

Next consider an antihole $\bar{C}_k = \{\boldsymbol{x}_0, \boldsymbol{x}_1, \ldots, \boldsymbol{x}_{k-1}\}$, listed in order, with $k \geq 6$. The neighbours in $\bar{C}_k$ of $\boldsymbol{x}_i$, for any $0 \leq i \leq k-1$, are $\Gamma_i = \{\boldsymbol{x}_{i+2}, \boldsymbol{x}_{i+3}, \ldots, \boldsymbol{x}_{i-2}\}$. We claim that either every vertex of Γ_i is in $\boldsymbol{x}_i^{\nwarrow}$ or every vertex of Γ_i is in $\boldsymbol{x}_i^{\searrow}$. To see this note that $(\boldsymbol{x}_{i+2}, \boldsymbol{x}_{i+3})$ is not an edge, and therefore $\{\boldsymbol{x}_{i+2}, \boldsymbol{x}_i, \boldsymbol{x}_{i+3}\}$ is an induced path in $G_{\succeq}$. By Lemma 3.3, without loss of generality, both $\boldsymbol{x}_{i+2}$ and $\boldsymbol{x}_{i+3}$ are in $\boldsymbol{x}_i^{\nwarrow}$. But $\{\boldsymbol{x}_{i+3}, \boldsymbol{x}_i, \boldsymbol{x}_{i+4}\}$ is also an induced path in $G_{\succeq}$. Consequently, as $\boldsymbol{x}_{i+3}$ is in $\boldsymbol{x}_i^{\nwarrow}$, Lemma 3.3 implies that $\boldsymbol{x}_{i+4}$ is in $\boldsymbol{x}_i^{\nwarrow}$. Repeating this argument through to the induced path $\{\boldsymbol{x}_{i-3}, \boldsymbol{x}_i, \boldsymbol{x}_{i-2}\}$ gives the claim.

Now consider the three vertices $\boldsymbol{x}_0, \boldsymbol{x}_2$ and $\boldsymbol{x}_4$. Since $k \geq 6$ these vertices are pairwise adjacent in $\bar{C}_k$. Without loss of generality, by the claim, Γ_0 is in $\boldsymbol{x}_0^{\nwarrow}$. Thus, $\boldsymbol{x}_2$ and $\boldsymbol{x}_4$ are in $\boldsymbol{x}_0^{\nwarrow}$. However $\boldsymbol{x}_0$ is in $\Gamma_2 \cap \Gamma_4$. Thus every vertex in Γ_2 is in $\boldsymbol{x}_2^{\searrow}$ and every vertex in Γ_2 is in $\boldsymbol{x}_4^{\searrow}$. Hence, $\boldsymbol{x}_4$ is in $\boldsymbol{x}_2^{\searrow}$ and $\boldsymbol{x}_2$ is in $\boldsymbol{x}_4^{\searrow}$, a contradiction. $\square$

Lemmas 3.4 and 3.5 together show, by applying the Strong Perfect Graph Theorem, that the auxiliary undirected graph is perfect.

Theorem 3.2. *The 2D auxiliary revealed preference graph $G_{\succeq}$ is perfect.* □

A Polynomial Time Algorithm in 2-Commodity Markets. In classical work, Grötschel, Lovász and Schrijver [20, 21] show that the VERTEX COVER problem in a perfect graph can be solved in polynomial time via the ellipsoid method.

Theorem 3.3. *[20] The* VERTEX COVER *problem is solvable in polynomial time on a perfect graph.* □

But by Theorem 3.2, the auxiliary undirected graph is perfect. Since the consumer rationality problem for two commodities corresponds to a vertex cover problem on this auxiliary undirected graph, we have:

Theorem 3.4. *In a two-commodity market, the* CONSUMER RATIONALITY *problem is solvable in polynomial time.* □

A Combinatorial Algorithm in 2-Commodity Markets. In his external report on the initial submission of this thesis, Sergey Norin pointed out that one can show these auxiliary graphs are actually *comparability graphs* using only Lemma 3.3. A comparability graph is a graph obtained from a partially ordered set P as follows: Let P be a set of elements, and $<$ be a relation on the elements of P which is acyclic and transitive. Construct a graph G on the ground set P by connecting any elements which are comparable.

Theorem 3.5. *The 2D auxiliary revealed preference graph $G_{\succeq}$ is a comparability graph.*

Proof. We will construct an orientation of the edges of the auxiliary graph, and then show that this orientation represents a partial order on the vertices.

Let $\{\boldsymbol{x}_1, \boldsymbol{x}_2\}$ be an edge of $G_{\succeq}$. As argued above, we either have $\boldsymbol{x}_1 \in \boldsymbol{x}_2^{\searrow}$ or $\boldsymbol{x}_2 \in \boldsymbol{x}_1^{\searrow}$, but not both, as we have assumed that points are distinct. Orient the edge from $\boldsymbol{x}_1$ to $\boldsymbol{x}_2$ if $\boldsymbol{x}_1 \in \boldsymbol{x}_2^{\searrow}$. We claim that repeating this for every edge constructs a transitive relation.

Consider any $\boldsymbol{x}_1$, $\boldsymbol{x}_2$, $\boldsymbol{x}_3$ for which we have constructed arcs $(\boldsymbol{x}_1, \boldsymbol{x}_2)$ and $(\boldsymbol{x}_2, \boldsymbol{x}_3)$. This would imply that $G_{\succeq}$ contained the edges $\{\boldsymbol{x}_1, \boldsymbol{x}_2\}$ and $\{\boldsymbol{x}_2, \boldsymbol{x}_3\}$. However, we have $\boldsymbol{x}_1 \in \boldsymbol{x}_2^{\searrow}$ and $\boldsymbol{x}_3 \in \boldsymbol{x}_2^{\nwarrow}$. By Lemma 3.3, this implies that the auxiliary graph also contained an $\{\boldsymbol{x}_1, \boldsymbol{x}_3\}$ edge, and it must be oriented $(\boldsymbol{x}_1, \boldsymbol{x}_3)$. Thus, the orientation of the edges forms a transitive relation, as desired. □

It is known that comparability graphs are perfect, which would give us our desired result. However, it is also known that the VERTEX COVER problem on comparability graphs of n vertices reduces to a network-flow problem on a network of $2n+2$ nodes. Both of these results are outlined in [32]. The network-flow reduction provides a straightforward combinatorial algorithm which is much simpler to implement than the ellipsoid-method algorithm provided in Theorem 3.3.

§3.4 3-Commodity Markets and Oriented Disk Graphs

We have shown that for two commodities, the consumer rationality problem can be solved in polynomial time. We now prove the problem is NP-complete if there are three (or more) commodities by presenting a reduction from PLANAR 3-SAT. The proof has three parts: first we transform an instance of PLANAR 3-SAT to an instance of VERTEX COVER in an associated undirected *gadget graph*. Second, we show that a vertex cover in the gadget graph corresponds to a directed feedback vertex set in a directed *oriented disc graph*. Finally, we prove that every oriented disc graph corresponds to a preference graph in a three-commodity market. Consequently, we can solve PLANAR 3-SAT using an algorithm for the three-commodity case of the CONSUMER RATIONALITY problem.

We begin by defining the class of oriented-disc graphs. Let $\{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n\}$ be points in the plane and let $\{B_1, \ldots, B_n\}$ be closed discs of varying radii such that B_i contains $\boldsymbol{x}_i$ on its boundary. We call this collection of points and discs an *oriented-disc drawing.* Given a drawing, we construct a directed graph $D = (V, A)$ on the vertex set $V = \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n\}$. There is an arc from $\boldsymbol{x}_i$ to $\boldsymbol{x}_j$ in D if $\boldsymbol{x}_j$, $j \neq i$, is contained in the disc B_i. A directed graph that can be built in this manner is called an *oriented-disc graph.*

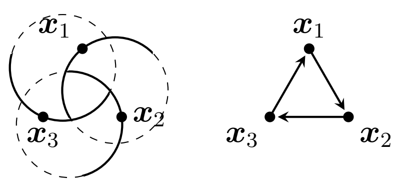

Figure 3.4.1: An oriented disc drawing and its corresponding oriented disc graph.

An example is given in Figure 3.4.1. The oriented-disc drawing is shown on the left and the the resulting oriented disc graph, a directed cycle on 3 vertices, is shown on the right. (We remark that, for enhanced clarity in the larger figures that follow, the boundary circles are drawn half-dotted.) Note that, even if the discs have uniform radii, the resulting oriented-disc graphs need not be symmetric – that is, $(\boldsymbol{x}_i, \boldsymbol{x}_j)$ can be an arc even if $(\boldsymbol{x}_j, \boldsymbol{x}_i)$ is not. This is due to the fact that $\boldsymbol{x}_i$ lies on the boundary, not at the centre, of its disc B_i. We now start by proving the third part of the reduction: every oriented disc graph corresponds to a preference graph in a three-commodity market.

Lemma 3.6. *Every oriented-disc graph corresponds to a preference graph induced by consumer data in a three-commodity market.*

Proof. Let D be any oriented-disc graph. We wish to build a three-commodity data set whose preference graph is D. Recall that the plane is homomorphic to the 2-dimensional sphere minus a point. Moreover, the inverse of the *stereographic projection* is a map from the plane to a sphere which preserves the shape of circles; see, for example, [14]. This motivates us to attempt to draw the points and discs on the unit sphere centered at $(1, 1, 1) \in \mathbf{R}^3$. To do this, we scale the oriented-disc drawing appropriately and embed it in a small region on the "underside" of the sphere, that is, around the point where the inwards normal vector is $(1, 1, 1)$. An example of this, where the oriented-disc graph is the directed 3-cycle, is shown in Figure 3.4.2(a).

We now need to create the corresponding collection of consumer data. Let $\{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n\}$ be the n points of some oriented-disc drawing of D embedded onto the underside of the sphere. Note that the intersection of a sphere and a plane is a circle. Furthermore, a plane through a point on the sphere will create a circle containing that point. Thus we may select the $\boldsymbol{x}_i$ to be the bundles chosen by the market and we may choose $\boldsymbol{p}_i$ such that the plane with normal $\boldsymbol{p}_i$ that passes through $\boldsymbol{x}_i$ intersects the sphere exactly along the boundary of the embedding of the disc B_i. An example is shown in Figure 3.4.2(b). Because $\boldsymbol{p}_i$ is non-negative it points into the sphere. Therefore, $\boldsymbol{x}_i$ is revealed preferred to every point on the inside of the embedding of B_i; it is not revealed preferred to any other point on the sphere. Hence, the preference graph $D_{\succeq}$ is isomorphic to the original oriented-disc graph, as desired. □

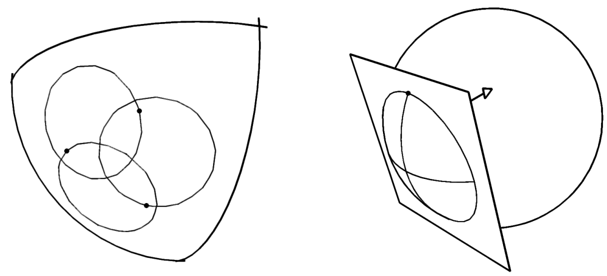

(a) (b)

Figure 3.4.2: A 3-cycle embedded on a sphere section, and a disc on a sphere.

Now, recall the first part of the reduction: we wish to transform an instance of PLANAR 3-SAT to an instance of VERTEX COVER in an associated undirected *gadget graph*. Our gadget graph is based upon a network used by Wang and Kuo [44] to prove the hardness of MAXIMUM INDEPENDENT SET in undirected unit-disc graphs. However, we are able to simplify their non-planar network by using an instance of PLANAR 3-SAT rather than the general 3-SAT. This simplification will be useful when implementing the second part of the reduction.

Let φ be an instance of PLANAR 3-SAT with variables $u_1, \ldots, u_n$ and clauses $C_1, \ldots, C_m$. Recall that φ is *planar* if the bipartite graph H_φ consisting of a vertex for each variable, a vertex for each clause, and edges connecting each clause to its three variables, is planar. The associated, undirected, *gadget graph* G_φ is constructed as shown in Figure 3.4.3. For each clause $C = (u_i \vee u_j \vee u_k)$, add a 3-cycle to the graph whose vertices are labelled by the appropriate literals for the variables u_i, u_j and u_k. We call these the *clause gadgets*. For each variable u_i, add a large cycle of even length whose vertices are alternatingly labelled as the literals u_i and $\bar{u}_i$. We call these the *variable gadgets*. Finally, add an edge from each variable in the clause gadgets to some vertex on the corresponding variable gadget with the opposite label – we choose a different variable vertex for each clause it is contained in.

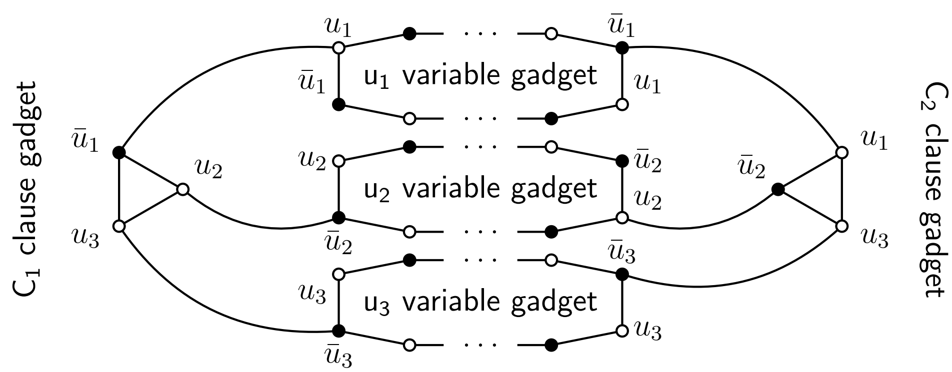

Figure 3.4.3: The gadget graph G_φ for $\varphi = (\bar{u}_1 \vee u_2 \vee u_3) \wedge (u_1 \vee \bar{u}_2 \vee u_3)$.

The next lemma is equivalent to the result shown by Wang and Kuo [44].

Lemma 3.7. *[44] The* PLANAR 3-SAT *instance φ is satisfiable if and only if G_φ has vertex cover set of size at most $2m + \frac{1}{2}\sum_{i=1}^{n} r_i$, where r_i is the number of vertices in the variable gadget's cycle for u_i.*

Proof. Suppose φ is satisfiable. Take any satisfying assignment, and let U be the set of literals which take TRUE values in the assignment, *i.e.* the literal "u_i" if the variable u_i was assigned TRUE, and the literal "$\bar{u}_i$" if the variable u_i was assigned FALSE. Let $\bar{U}$ be the remaining literals. Now, every vertex in a variable gadget of G_φ whose label is in $\bar{U}$ will be selected to be in the vertex cover. In total this amounts to $\frac{1}{2}\sum_{i=1}^{n} r_i$ vertices, and these *cover* every edge in the variable gadgets of G_φ. Next consider the clause gadgets of G_φ. We must select two vertices of each clause gadget to cover the edges of the 3-cycle. This amounts to $2m$ vertices. Since we have a satisfying assignment and we chose the nodes corresponding to $\bar{U}$ in the variable gadgets, each clause gadget must have at least one incident edge covered by the variable gadgets' selected vertices. Hence, selecting the other two vertices will cover all incident edges to the clause gadgets, and all edges of the gadget's cycle. Thus we have a vertex cover with $2m + \frac{1}{2}\sum_{i=1}^{n} r_i$ nodes. For example, in Figure 3.4.3, if we set all variables to FALSE, one possible vertex cover is the set of vertices labelled with non-negated literals, *i.e.* those coloured in white. (Note, clearly, the set of white vertices will not typically form a vertex cover in the gadget graph.)

Conversely, suppose we have a vertex cover $\mathcal{C}$ containing at most $2m+\sum_{i=1}^{n} \frac{r_i}{2}$ vertices. Each variable gadget must contribute at least $\frac{r_i}{2}$ vertices, otherwise we cannot cover every edge in its cycle. Each clause gadget must contribute at least two vertices, or one edge in the 3-cycle will be uncovered. Hence, $\mathcal{C}$ contains exactly $2m + \sum_{i=1}^{n} \frac{r_i}{2}$ vertices. The $\frac{1}{2}r_i$ vertices from the variable gadget for u_i corresponds either to the set of all vertices with negated labels or to the set with non-negated labels, otherwise there is an uncovered edge in the cycle. This induces a truth assignment; set u_i to TRUE if all the "$\bar{u}_i$"-labelled vertices are selected, and FALSE if the "u_i"-labelled vertices are selected. Furthermore this is a satisfying assignment. To see this note that as $\mathcal{C}$ covers all edges, the unselected vertex in each clause is a literal which evaluates to TRUE by the selected assignment. □

Hence, to solve for the satisfiability of φ, it suffices to test whether G_φ admits a vertex cover with at most $2m + \frac{1}{2}\sum_{i=1}^{n} r_i$ vertices. It remains to show the second of the three parts of the reduction. That is, we need to show that this VERTEX COVER problem in the undirected gadget graph can be solved by finding a minimum directed feedback vertex set in an oriented-disc graph D_φ. The basic idea is straightforward (albeit that the implementation is intricate). The oriented-disc graph D_φ will contain a digon for each edge in some G_φ. However, it will also contain a collection of additional arcs. The key fact will be that these additional arcs form an acyclic subgraph of D_φ. Thus every cycle in D_φ must induce a digon. Consequently, a minimum directed feedback vertex set need only intersect each digon to ensure that every cycle is hit. As argued previously, hitting the underlying graph formed by the digons of D_φ corresponds to selecting a vertex cover in G_φ, as desired. We now formalise this argument.

Lemma 3.8. *For every instance φ of* PLANAR 3-SAT*, there exists an oriented-disc graph D_φ on which the* DIRECTED FEEDBACK VERTEX SET *problem is equivalent to the* VERTEX COVER *problem on G_φ.*

Proof. We prove this by explicitly constructing the oriented-disc drawing. Recall the disc graph D_φ should contain a digon for each edge in G_φ. To do this, we begin with sufficiently a large planar drawing of H_φ, the planar bipartite network associated with φ. At each clause vertex, we place an oriented-disc construction for the clause gadget. This construction, along with its resulting graph, is shown in Figure 3.4.4. The figure shows a clause gadget and a section of each of the neighbouring three variable gadgets to which it is attached. Observe from the figure that, as claimed, the set of arcs created in D_φ which are not in a digon, form an acyclic subgraph of D_φ.

It remains to construct the large cycles for the variable gadgets, and connect them to the clause gadgets. However, parts of these cycles are already included in the clause gadgets. Thus, it suffices to join these cycle segments together via paths of digons. This can be done via the oriented disc constructions shown in Figure 3.4.5. To draw the cycle for some variable, say u_i, we note that u_i's vertex in the planar network H_φ shares and edge with every clause gadget which connects to u_i's gadget. Hence, as illustrated in Figure 3.4.6, we may follow along the edges of H_φ to construct the cycle. For example, in the figure, the

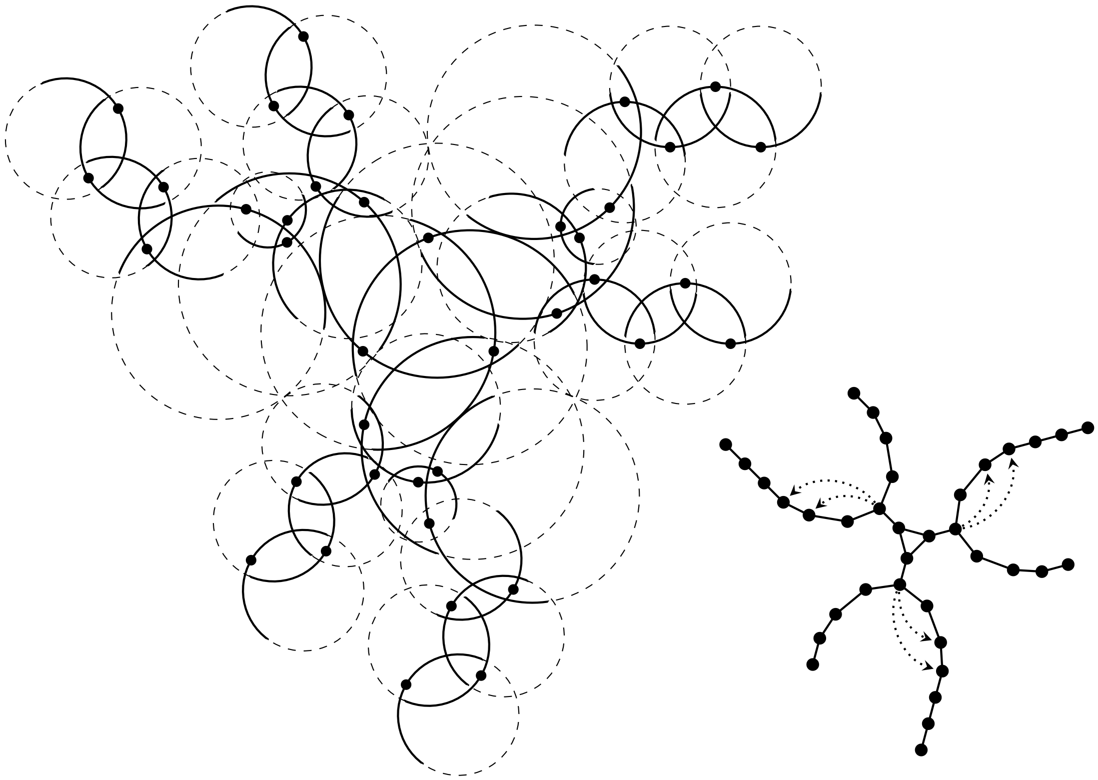

Figure 3.4.4: Oriented-disc construction of the clause gadget, and its resulting graph.

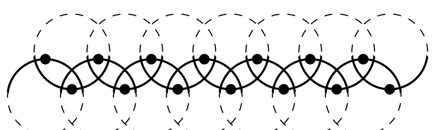

(a) Straight line

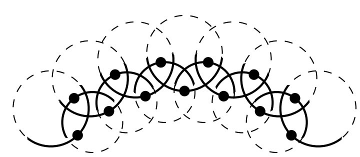

(b) Curve

Figure 3.4.5: Paths of bidirected edges as oriented-disc drawings.

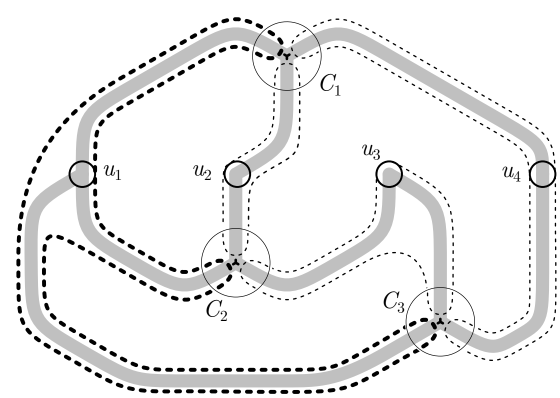

Figure 3.4.6: G_φ as an oriented-disc graph.

variable cycle for u_1 (highlighted) follows the topology of the edges incident to u_1's vertex, and joins the clause gadgets (circled) to one another.

Observe that constructions in Figure 3.4.5 produce paths of digons in D_φ, where every arc produced is contained in a digon. It follows that the only arcs in D_φ that are not in digons are in the neighbourhoods of the clause gadgets and, as we have seen, these are acyclic. But then, to hit all the cycles in D_φ, it suffices to hit all the digons, which, in turn, corresponds to a vertex cover in G_φ, completing the proof. □

This completes all the steps in the reduction and we obtain:

Theorem 3.6. *The* CONSUMER RATIONALITY *problem is NP-complete for a market with at least 3 commodities.* □

Chapter 4

Conclusion

This thesis aimed to study two relaxations of the axioms of revealed preference, one additive and one combinatorial. We have shown the additive rule is efficiently computable in the auction setting, which led us to study implementations of the rule, and variations thereof. Further research is necessary to determine whether these variations have their intended effects in empirical settings.

The combinatorial relaxation, however, is NP-hard in the standard setting (unless the market has only 2 commodities), and therefore it is unlikely to be a good candidate for an activity rule. We leave open the problem of finding the approximation complexity of this problem: a reasonable degree of approximation may turn out to be efficiently computable. Furthermore, the problem is NP-hard in the worst case, but it may not be too difficult to compute for real-world data sets and simulated auctions. Further study is required to determine whether such a rule may be worth considering in practice. Such implementation questions are, however, outside of the scope of this thesis.

In showing that the combinatorial problem was efficiently computable in a 2-commodity market, we showed that undirected auxiliary graphs form a class of perfect graphs. It is our belief that this is a previously unknown class of perfect graphs, and it may be interesting to characterise this class. This also leads to the problem of characterising the class of preference graphs feasible for a market of given dimension. Both these problems are left open.

References

[1] S. Afriat, “The construction of a utility function from expenditure data”, *International Economic Review*, **8**, pp67–77, 1967.

[2] S. Afriat, “On a system of inequalities in demand analysis: an extension of the classical method”, *International Economic Review*, **14**, pp460-472, 1967.

[3] G. Akerlof and J. Yellen, “Can small deviations from rationality make significant differences to economic equilibria?”, *Amer. Econ. Rev.*, **75(4)**, pp708–720, 1985.

[4] J. Apesteguia and M. Ballester, “A measure of rationality and welfare”, to appear in *Journal of Political Economy*, 2015.

[5] L. Ausubel and O. Baranov, “Market design and the evolution of the combinatorial clock auction”, *Amer. Econ. Rev.*, **104(5)**, pp446–451, 2014.

[6] L. Ausubel, P. Cramton and P. Milgrom, “The clock-proxy auction: a practical combinatorial auction design”, in P. Cramton, Y. Shoham and R. Steinberg (eds.), *Combinatorial Auctions*, pp115–138, MIT Press, 2006.

[7] C. Berge, “Färbung von Graphen deren sämtliche beziehungsweise deren ungerade Kreise starr sind (Zusammenfassung)”, *Wiss. Z. Martin-Luther-Univ. Halle-Wittenberg Math.-Natur. Reihe*, **10**, pp114-115, 1961.

[8] C. Chambers and F. Echenique, “Supermodularity and preferences”, *Journal of Economic Theory*, **144(3)**, pp1004–1014, 2009.

[9] M. Chudnovsky, N. Robertson, P. Seymour and R. Thomas. “The strong perfect graph theorem”, *Annals of Mathematics*, pp51–229, 2006.

[10] P. Cramton, “Spectrum auction design”, *Review of Industrial Organization*, **42(2)**, pp161–190, 2013.

[11] M. Dean and D. Martin, “Measuring rationality with the minimum cost of revealed preference violations”, to appear in *Review of Economics and Statistics*, 2015.

[12] R. Deb and M. Pai, “The geometry of revealed preference”, *Journal of Mathematical Economics*, **50**, pp203–207, 2014.

[13] I. Dinur and S. Safra, “On the hardness of approximating minimum vertex cover”, *Annals of Mathematics*, **162(1)**, pp439-485, 2005.

[14] Earl, R. (2007). Geometry II: 3.1 Stereographic Projection and the Riemann Sphere. Retrieved from https://people.maths.ox.ac.uk/earl/G2-lecture5.pdf

[15] F. Echenique, D. Golovin and A. Wierman, “A revealed preference approach to computational complexity in economics”, *Proceedings of EC*, pp101–110, 2011.

[16] F. Echenique, S. Lee and M. Shum, “The money pump as a measure of revealed preference violations”, *Journal of Political Economy*, **119(6)**, pp1201–1223, 2011.

[17] M. Famulari, “A household-based, nonparametric test of demand theory”, *Review of Economics and Statistics*, **77**, pp372-383, 1995.

[18] G. Gigerenzer and R. Selten (eds), *Bounded Rationality: the Adaptive Toolbox*, MIT Press, 2001.

[19] J. Gross, "Testing data for consistency with revealed preference", *The Review of Economics and Statistics*, **77(4)**, pp701–710, 1995.

[20] M. Grötschel, L. Lovász and A. Schrijver, "Polynomial algorithms for perfect graphs", *Annals of Discrete Mathematics*, **21**, pp325-356, 1984.

[21] M. Grötschel, L. Lovász and A. Schrijver, *Geometric Algorithms and Combinatorial Optimisation*, Springer–Verlag, 1988.

[22] F. Gul and E. Stacchetti, "Walrasian equilibrium with gross substitutes", *Journal of Economic Theory*, **87**, pp95–124, 1999.

[23] V. Guruswami, J. Hastad, R. Manokaran, P. Raghavendra, and M. Charikar, "Beating the random ordering is hard: Every ordering CSP is approximation resistant", *SIAM J. Computing*, **40(3)**, pp878–914, 2011.

[24] P. Harsha, C. Barnhart, D. Parkes and H. Zhang, "Strong activity rules for iterative combinatorial auctions", *Computers and O.R.*, **37(7)**, pp1271–1284, 2010.

[25] J. Heufer, "A geometric approach to revealed preference via Hamiltonian cycles", *Theory and Decision*, **76(3)**, pp329–341, 2014.

[26] H. Houthakker, "Revealed Preference and the Utility Function", *Economica, New Series*, **17(66)**, pp159–17, 1950.

[27] M. Houtman and J. Maks, "Determining all maximal data subsets consistent with revealed preference", *Kwantitatieve Methoden*, **19**, pp89–104, 1950.

[28] R. Karp, “Reducibility among combinatorial problems”, *Complexity of Computer Computations*, New York: Plenum, pp85–103, 1972.

[29] R. Karp, “A characterization of the minimum cycle mean in a digraph”, *Discrete Mathematics*, **23(3)**, p.309–311, 1978.

[30] S. Khot, “On the power of unique 2-prover 1-round games”, Proceedings of STOC, pp767–775, 2002.

[31] L. Lovász, “Normal hypergraphs and the perfect graph conjecture”, *discrete Mathematics*, **2(3)**, pp253–267, 1972.

[32] R. Möhring., “Algorithmic Aspects of Comparability Graphs and Interval Graphs”, *Graphs and Order*, pp. 41–101, D. Reidel Publishing Company, 1985

[33] H. Rose. “Consistency of preference: the two-commodity case”, *Review of Economic Studies*, **25**, pp124–125, 1958.

[34] P. Samuelson, “A note on the pure theory of consumer’s behavior”, *Economica*, **5(17)**, pp61–71, 1938.

[35] P. Samuelson, “Consumption theory in terms of revealed preference”, *Economica*, **15(60)**, pp243–253, 1948.

[36] P. Seymour, “Packing directed circuits fractionally”, *Combinatorica*, **15(2)**, pp281–288, 1995.

[37] O. Svensson, “Hardness of vertex deletion and project scheduling”, In Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, volume 7408 of Lecture Notes in Computer Science, pages 301–312. 2012.

[38] J. Swafford and G. Whitney, "Nonparametric test of utility maximization and weak separability for consumption, leisure and money", *Review of Economics and Statistics*, **69**, pp458-464, 1987.

[39] H. Varian, "The nonparametric approach to demand analysis", *Econometrica*, **50(4)**, pp945-973, 1982.

[40] H. Varian, "Goodness-of-fit in optimizing models", *Journal of Econometrics*, **46**, pp125-140, 1990.

[41] H. Varian, "Revealed preference", in Szenberg, Ramrattand and Gottesman (eds.), *Samulesonian Economics and the* 21*st Century*, pp99–115, O.U.P., 2005.

[42] H. Varian, "Position auctions", *Int. J. Ind. Organ.*, **25(6)**, pp1163–1178, 2007.

[43] R. Vohra, *Mechanism Design: A Linear Programming Approach*, C.U.P., 2011.

[44] D. Wang and Y. Kuo, "A study on two geometric location problems", *Information Processing Letters*, **28(6)**, pp281–286, 1988.