

The Combinatorial World (of Auctions) According to GARP

Shant Boodaghians¹ and Adrian Vetta²

¹ Department of Mathematics and Statistics, McGill University
shant.boodaghians@mail.mcgill.ca

² Department of Mathematics and Statistics, and School of Computer Science, McGill University
vetta@math.mcgill.ca

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 applications such as combinatorial auctions. In the auction setting, we present an efficient combinatorial algorithm to find a virtual valuation function with the optimal (additive) rationality guarantee. Moreover, we show that there exists such a valuation function that both is individually rational and is minimum (that is, it is component-wise dominated by any other individually rational, virtual valuation function that approximately fits the data). Similarly, given upper bound constraints on the valuation function, we show how to fit the maximum virtual valuation function with the optimal additive rationality guarantee. In practice, revealed preference bidding constraints are very demanding. We explain how approximate rationality can be used to create relaxed revealed preference constraints in an auction. We then 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.

1 Introduction

Underlying the theory of consumer demand is a standard rationality assumption: given a set of items with price vector \mathbf{p} , a consumer will demand the bundle \mathbf{x} of maximum utility whose cost is at most her budget B . Of fundamental import, therefore, is whether or not the decision making behaviour of a real consumer is consistent with the maximization of a utility function. Samuelson [18,19] introduced *revealed preference* to provide a theoretical framework within which to analyse this question. Furthermore, this concept now lies at the heart of current empirical work in the field; see, for example, Gross [11] and Varian [23]. Specifically, Samuelson [18] conjectured that the *weak axiom of revealed preference* (WARP) was a necessary and sufficient condition for *integrability* – the ability to construct a utility function which fits observed behaviour.

However, Houtthakker [14] proved that the weak axiom was insufficient. Instead, he presented a *strong axiom of revealed preference* (SARP) and showed non-constructively that it was necessary and sufficient in the case where behaviour is determined via a single-valued demand function. Afriat [1] provided an extension to multi-valued demand functions – where ties are allowed – by showing that the *generalized axiom of revealed preference* (GARP) is necessary and sufficient for integrability.¹ Furthermore, Afriat’s approach was constructive (producing monotonic, concave, piecewise-linear utility functions) and applied to the setting of a finite collection of observational data. This rendered his method more suitable for practical use.

In addition to its prominence in testing for rational behaviour, revealed preference has become an important tool in mechanism design. A notable area of application is auction design. For combinatorial auctions, Ausubel, Cramton and Milgrom [4] proposed bidding activity rules based upon WARP. 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 bidding rules have proved successful because they are extremely difficult to game [6]. Harsha et al. [13] examine GARP-based bidding rules, and Ausubel and Baranov [3] 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.

Revealed preference also plays a key role in motivating the *generalised second price mechanism* used in adword auctions. Indeed, these position auctions have welfare maximizing solutions with respect to a revealed preference equilibrium concept; see [23] and [8].

1.1 Our Results

Multiple methods have been proposed to approximately measure how consistent a data set is with rational behaviour; see Gross [11] for a comparison of a sample of these approaches. In this paper, 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 3 that an individually rational virtual valuation function can be obtained such that its additive deviation from rationality is exactly the minimum mean length

¹ Afriat [1] gave several equivalent necessary and sufficient conditions for integrability. One of these, *cyclical consistency*, is equivalent to GARP as shown by Varian [21].

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 [6] 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 4, 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/best-case welfare guarantees.

2 Revealed Preference

In this section, we first review revealed preference. We then examine its use in auction design and describe how to formulate it in terms of a bidding graph.

2.1 Revealed Preference with Budgets

The standard revealed preference model instigated by Samuelson [18] is as follows. We are given a set of observations

$$\{(B_1, \mathbf{p}_1, \mathbf{x}_1), (B_2, \mathbf{p}_2, \mathbf{x}_2), \dots, (B_T, \mathbf{p}_T, \mathbf{x}_T)\} .$$

At time t , $1 \leq t \leq T$, the set of items has a price vector \mathbf{p}_t and the consumer chooses to spend her budget B_t on the bundle \mathbf{x}_t .² We say that bundle \mathbf{x}_t is (directly) revealed preferred to bundle \mathbf{y} , denoted $\mathbf{x}_t \succeq \mathbf{y}$, if \mathbf{y} was affordable when \mathbf{x}_t was purchased. We say that bundle \mathbf{x}_t is strictly revealed preferred to bundle \mathbf{y} , denoted $\mathbf{x}_t \succ \mathbf{y}$, if \mathbf{y} was (strictly) cheaper than \mathbf{x}_t when \mathbf{x}_t was purchased. This gives *revealed preference* (1) and *strict revealed*

² It is not necessary to present the model in terms of “time”. We do so because this best accords with the combinatorial auction application.

preference (2):

$$\mathbf{p}_t \cdot \mathbf{y} \leq \mathbf{p}_t \cdot \mathbf{x}_t \Rightarrow \mathbf{x}_t \succeq \mathbf{y} \quad (1)$$

$$\mathbf{p}_t \cdot \mathbf{y} < \mathbf{p}_t \cdot \mathbf{x}_t \Rightarrow \mathbf{x}_t \succ \mathbf{y} \quad (2)$$

Furthermore, a basic assumption is that the consumer optimises a locally non-satiated utility function.³ Consequently, at time t she will spend her entire budget, *i.e.*, $\mathbf{p}_t \cdot \mathbf{x}_t = B_t$. In the absence of ties, preference orderings give relations that are anti-symmetric and transitive. This leads to an axiomatic approach to revealed preference formulated in terms of WARP and SARP by Houthakker [14]. The *weak axiom of revealed preference* (WARP) states that the relation should be asymmetric, *i.e.* $\mathbf{x} \succeq \mathbf{y} \Rightarrow \mathbf{y} \not\succeq \mathbf{x}$. Its transitive closure, the *strong axiom of revealed preference* (SARP) states that the relation should be acyclic. Our interest lies in the general case where ties are allowed. This produces what we dub the *k-th Axiom of Revealed Preference* (KARP): Given a fixed integer k and any $\kappa \leq k$

$$\mathbf{x}_t = \mathbf{x}_{t_0} \succ \mathbf{x}_{t_1} \succ \cdots \succ \mathbf{x}_{t_{\kappa-1}} \succ \mathbf{x}_{t_\kappa} = \mathbf{y} \Rightarrow \mathbf{y} \not\succeq \mathbf{x}_t . \quad (3)$$

There are two very important special cases of KARP. For $k = 1$, this is simply WARP, *i.e.* $\mathbf{x}_t \succeq \mathbf{y} \Rightarrow \mathbf{y} \not\succeq \mathbf{x}_t$. This is just the basic property that for a preference ordering, we cannot have that \mathbf{y} is strictly preferred to \mathbf{x}_t if \mathbf{x}_t is preferred to \mathbf{y} . On the other hand, suppose we take k to be arbitrarily large (or simply larger than the total number of observed bundles). Then we have the *Generalized Axiom of Revealed Preference* (GARP), the simultaneous application of KARP for each value of k . In particular, GARP encodes the property of transitivity of preference relations. Specifically, for any k , if

$$\mathbf{x}_t = \mathbf{x}_{t_0} \succeq \mathbf{x}_{t_1} \succeq \cdots \succeq \mathbf{x}_{t_{k-1}} \succeq \mathbf{x}_{t_k} = \mathbf{y} ,$$

then, by transitivity, $\mathbf{x}_t \succeq \mathbf{y}$. The first axiom of revealed preference then implies that $\mathbf{y} \not\succeq \mathbf{x}_t$.

The underlying importance of GARP follows from a classical result of Afriat [1]: there exists a nonsatiated, monotone, concave utility function that rationalizes the data if and only if the data satisfy GARP. Brown and Echenique [5] examine the setting of indivisible goods and Echenique et al. [7] consider the consequent computational implications.

³ Local non-satiation states that for any bundle \mathbf{x} there is a more preferred bundle arbitrarily close to \mathbf{x} . A monotonic utility function is locally non-satiated, but the converse need not hold.

2.2 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 2.1. 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 from perfect capital markets. (This assumption is slightly unrealistic; Harsha et al. [13] show how to implement a budgeted revealed preference model for combinatorial auctions; see also Section 4.4).

Third, the observations $(\mathbf{p}_t, \mathbf{x}_t)$, 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 \mathbf{x}_t .

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

$$v(\mathbf{x}_t) - \mathbf{p}_t \cdot \mathbf{x}_t \geq v(\mathbf{x}_s) - \mathbf{p}_t \cdot \mathbf{x}_s \quad (4)$$

$$v(\mathbf{x}_s) - \mathbf{p}_s \cdot \mathbf{x}_s \geq v(\mathbf{x}_t) - \mathbf{p}_s \cdot \mathbf{x}_t \quad (5)$$

Summing Inequalities (4) and (5) and rearranging gives

$$(\mathbf{p}_t - \mathbf{p}_s) \cdot \mathbf{x}_s \geq (\mathbf{p}_t - \mathbf{p}_s) \cdot \mathbf{x}_t \quad (6)$$

This is the revealed preference condition for combinatorial auctions proposed as a bidding activity rule by Ausubel, Crampton and Milgrom [4]. The activity rule simply states that, between time s and time t , the price of bundle \mathbf{x}_s must have risen by at least as much as the price of \mathbf{x}_s . If condition (6) is not satisfied then the auction mechanism will not allow the later bid to be made.

Observe that the bidding rule (6) 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 $\mathbf{x} = (x_1, \dots, x_n)$ and an amount x_0 of money we denote by

$\hat{\mathbf{x}} = (x_0, x_1, \dots, x_n)$ the concatenation of x_0 and \mathbf{x} . If $\mathbf{p} = (p_1, \dots, p_n)$ is the price vector for the non-monetary items, then $\hat{\mathbf{p}} = (1, p_1, \dots, p_n)$ gives the prices of all items including money.

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

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

Equivalently,

$$v(\mathbf{x}_t) \geq v(\mathbf{x}_s) + \mathbf{p}_t \cdot \mathbf{x}_t - \mathbf{p}_t \cdot \mathbf{x}_s . \quad (7)$$

But Inequality (7) is equivalent to Inequality (4). Inequality (5) follows symmetrically, and together these give the revealed preference bidding rule (6). 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 \mathbf{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

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

This induces the inequality

$$v(\mathbf{x}_i) - \mathbf{p}_i \cdot \mathbf{x}_i \geq v(\mathbf{x}_{i+1}) - \mathbf{p}_i \cdot \mathbf{x}_{i+1} . \quad (8)$$

Summing (8) over all i , we obtain

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

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

$$(\mathbf{p}_k - \mathbf{p}_0) \cdot \mathbf{x}_0 \geq \sum_{i=1}^k (\mathbf{p}_i - \mathbf{p}_{i-1}) \cdot \mathbf{x}_i . \quad (9)$$

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 these rules.

2.3 A Graphical View of Revealed Preference

Given the set of price-bid pairings $\{(\mathbf{p}_t, \mathbf{x}_t) : 1 \leq t \leq T\}$, 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 \mathbf{x}_t , $1 \leq t \leq T$, there is an arc $(\mathbf{x}_t, \mathbf{y})$ for each bundle $\mathbf{y} \in V$. In order to define the length $\ell_{\mathbf{x}_t, \mathbf{y}}$ of an arc $(\mathbf{x}_t, \mathbf{y})$, note that Inequality (4) applied to $\mathbf{x}_s = \mathbf{y}$ gives

$$v(\mathbf{y}) \leq v(\mathbf{x}_t) + \mathbf{p}_t \cdot (\mathbf{y} - \mathbf{x}_t) ,$$

otherwise we would prefer bundle \mathbf{y} at time t . For the arc length, we would like to simply set $\ell_{\mathbf{x}_t, \mathbf{y}} = \mathbf{p}_t \cdot (\mathbf{y} - \mathbf{x}_t)$. Observe, however, that the bundle \mathbf{x}_t may be chosen in more than one time period. That is, possibly $\mathbf{x}_t = \mathbf{x}_{t'}$ for some $t \neq t'$. 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

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

Now the WARP-based bidding rule (6) of Ausubel et al. [4] is equivalent to

$$(\mathbf{p}_t - \mathbf{p}_s) \cdot \mathbf{x}_s - (\mathbf{p}_t - \mathbf{p}_s) \cdot \mathbf{x}_t \geq 0 .$$

However,

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

It is then easy to see that the bidding constraint (6) is violated if and only if the bidding graph contains no negative digons (cycles of length two). Furthermore, we can interpret KARP and GARP in 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 [25]. 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 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 $\mathbf{y} \notin \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T\}$ has zero out-degree in G . Consequently, \mathbf{y} cannot be contained in any cycle. Thus, it will suffice to consider only the subgraph induced by the bids $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T\}$. In a combinatorial auction there is typically one bid per time period and the number of periods is quite small.⁴ 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. [13]. 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 $O(T^3)$ 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.

3 Approximate Virtual Valuation Functions

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

Theorem 3.1. *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 3.2 below; see also [25]. From an economic perspective, however, what is most important is not whether agents are perfectly rational but “whether optimization is a reasonable way to de-

⁴ For example, in a bandwidth auction there are at most a few hundred rounds.

scribe some behavior” [22].⁵ It is then important to study the consequences of approximately rational behaviour, see, for example, Akerlof and Yellen [2]. First, though, is it possible to quantify the degree to which agents are rational? Gross [11] examines assorted methods to test the degree of rationality. Notable amongst them is the *Afriat Efficiency Index* [1,22]. Here the condition required to imply a preference is strengthened multiplicatively. Specifically, $\mathbf{x}_t \succeq \mathbf{y}$ only if $\mathbf{p}_t \cdot \mathbf{y} \leq \lambda \cdot \mathbf{p}_t \cdot \mathbf{x}_t$ where $\lambda < 1$. We examine this index with respect to the bidding graph in Section 4.4. For combinatorial auctions, a variant of this constraint was examined experimentally by Harsha et al. [13].

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 \mathbf{y} ,

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

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*.

3.1 Minimum Mean Cycles and Approximate Virtual Valuations

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 [15].

Theorem 3.2. *An ϵ -approximate valuation function which (approximately) rationalises bidding behaviour exists if and only if the bidding graph has minimum mean cycle $\mu(G) \geq -\epsilon$.*

⁵ 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, [10].

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

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

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 $(\mathbf{x}_t, \mathbf{y})$, we have

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

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

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

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

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

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

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

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

Recall that, the bidding behaviour is irrational only if $\mu(G)$ is strictly negative. We emphasize that Theorem 3.2 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 \mathbf{y} , that $\hat{v}(\mathbf{x}_t) - \mathbf{p}_t \cdot \mathbf{x}_t \geq \hat{v}(\mathbf{y}) - \mathbf{p}_t \cdot \mathbf{y} + \delta$. Thus, \mathbf{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.

3.2 Individually Rational Virtual Valuation Functions

Theorem 3.2 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 $\hat{v}(\mathbf{x}_t) - \mathbf{p}_t \cdot \mathbf{x}_t \geq 0$. But individual rationality is (almost certainly) violated for the the root node r since we have $\hat{v}(\mathbf{x}_r) = 0$.

It is possible to obtain an individually rational, approximate, virtual valuation function simply by taking the \hat{v} from Theorem 3.2 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.

The Minimum Individually Rational Virtual Valuation Function.

We say that $v(\cdot)$ is the *minimum individually rational, ϵ -approximate virtual valuation function* if $v(\mathbf{x}_t) \leq \omega(\mathbf{x}_t)$ for each $1 \leq t \leq T$, for any other individually rational, ϵ -approximate virtual valuation function $\omega(\cdot)$. 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 3.3. *The minimum individually rational, μ -approximate virtual valuation function exists and can be found in polynomial time.*

Proof. We create an auxiliary directed graph H from \hat{G} by adding a sink vertex \mathbf{z} . We add an arc $(\mathbf{x}_t, \mathbf{z})$ of length $-\mathbf{p}_t \cdot \mathbf{x}_t$, 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 $\hat{d}(\cdot)$ the shortest path distance from vertex \mathbf{x}_t to \mathbf{z} in H . We claim that setting $v(\mathbf{x}_t) = -\hat{d}(\mathbf{x}_t)$ gives the minimum individually rational, μ -approximate virtual valuation function.

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

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

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

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

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

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

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

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

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

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

But then

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

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

The Maximum (Individually Rational) 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 4. 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 $\mathbf{p}_{t+1} \cdot \mathbf{x}_t$ is an upper bound on the value of bundle \mathbf{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 $v(\mathbf{x}_i) \leq \beta_i$ for each $i \in I$. Then there is a *unique* maximum μ -approximate virtual valuation function.

Theorem 3.4. *Given a set of constraints, the maximum μ -approximate virtual valuation function exists and can be found in polynomial time.*

Proof. Let $v(\mathbf{x}_i) \leq \beta_i$ for each $i \in I$. We construct a graph H from \hat{G} by adding a source vertex \mathbf{z} with arcs of length β_i from \mathbf{z} to \mathbf{x}_i , for each $i \in I$. Since \mathbf{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* \mathbf{z} . We claim that setting $v(\mathbf{x}) = \hat{d}(\mathbf{x})$ 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 \mathbf{z} to \mathbf{x}_i . Thus the shortest path to \mathbf{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 $v(\mathbf{x}_t) \geq \omega(\mathbf{x}_t)$ for every bundle \mathbf{x}_t . For a contradiction, suppose that $P = \{\mathbf{z}, \mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_r\}$ is the shortest path from \mathbf{z} to \mathbf{y}_r in H and that $v(\mathbf{y}_r) < \omega(\mathbf{y}_r)$. Observe that the node adjacent to \mathbf{z} on P must be $\mathbf{y}_1 = \mathbf{x}_i$ for some $i \in I$. Now because $\omega()$ is a μ -approximate valuation function, we have

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

Cancelling terms produces

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

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

Notice that Theorem 3.4 does not guarantee that the maximum virtual valuation function is individually rational. For example, suppose $\beta_t = \mathbf{p}_t \cdot \mathbf{x}_t$, for all $1 \leq t \leq T$. Individual rationality then implies that $v(\mathbf{x}_t)$ must equal $\mathbf{p}_t \cdot \mathbf{x}_t$ 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 3.4 must be individually rational by maximality.

4 Revealed Preference Auction Bidding Rules

So far, we have focused upon how to test the degree of rationality reflected in a data set. Specifically, we saw in Theorem 3.2 that the minimum mean length of a cycle, $\mu(G)$, gives an exact and optimal goodness of fit measure for rationality. Furthermore, Theorem 3.3 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 tool in mechanism design. In particular, we saw in Section 2.2 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.

4.1 Relaxed Revealed Preference Bidding Rules

Consider a combinatorial auction at time (round) t where our prior price-bundle bidding pairs are $\{(\mathbf{p}_1, \mathbf{x}_1), (\mathbf{p}_2, \mathbf{x}_2), \dots, (\mathbf{p}_{t-1}, \mathbf{x}_{t-1})\}$. By Inequality (6) in section 2.2, rational bidding at time t implies that

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

Moreover, a necessary condition is then that $(\mathbf{p}_t - \mathbf{p}_s) \cdot \mathbf{x}_s \geq (\mathbf{p}_t - \mathbf{p}_s) \cdot \mathbf{x}_t$ 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 $(\mathbf{p}_t, \mathbf{x}_t)$ 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 ,

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

A necessary condition then is $(\mathbf{p}_t - \mathbf{p}_s) \cdot \mathbf{x}_s \geq (\mathbf{p}_t - \mathbf{p}_s) \cdot \mathbf{x}_t - 2\epsilon$, 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

$$(\mathbf{p}_k - \mathbf{p}_0) \cdot \mathbf{x}_0 \geq \sum_{i=1}^k (\mathbf{p}_i - \mathbf{p}_{i-1}) \cdot \mathbf{x}_i - (k+1) \cdot \epsilon \quad (10)$$

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 4.1. *A set of price-bid pairings $\{(\mathbf{p}_t, \mathbf{x}_t) : 1 \leq t \leq T\}$ 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 3.2, 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 $\{\mathbf{x}_i\}_{i=1}^k$ of bundles. Let t_i be the time when $\ell_{\mathbf{x}_i, \mathbf{x}_{i+1}}$ was minimized, and let $\mathbf{p}_i := \mathbf{p}_{t_i}$. Then we have

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

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. \square

4.2 Relaxed KARP-Based Bidding Rules

Theorem 4.1 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 4.2. *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 $\{e_1, e_2, \dots, e_{|C|}\}$ in order. Then

$$\sum_{i=1}^{|C|} \sum_{j=i}^{i+k-1} \ell_{e_j} = k \cdot \sum_{i=1}^{|C|} \ell_{e_i} = k \cdot \ell(C) = k \cdot |C| \cdot \frac{\ell(C)}{|C|} . \quad (11)$$

Above, the inner summation is taken modulo $|C|$. On the other hand take any path segment $P = \{e_i, e_{i+1}, \dots, e_{i+k-1}\}$, where again the subscript summation is modulo $|C|$. Because the graph is complete and the maximum arc length is ℓ^{\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,

$$\sum_{i=1}^{|C|} \sum_{j=i}^{i+k-1} \ell_{e_j} \geq -|C| \cdot \ell^{\max} . \quad (12)$$

Combining Equalities (11) and Inequality (12) 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 $-\frac{\ell^{\max}}{k}$. \square

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 4.1. *Given a set of price-bid pairings $\{(\mathbf{p}_t, \mathbf{x}_t) : 1 \leq t \leq T\}$ 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 (10) 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 $\ell'_{\mathbf{x}_s, \mathbf{x}_t} := \ell_{\mathbf{x}_s, \mathbf{x}_t} + \epsilon$. Then every cycle in G' of cardinality at most $k + 1$ has non-negative length. By Theorem 4.2, 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 3.2 and 3.3 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 4.1 can be improved. The answer is *no*; Theorem 4.2 is tight.

Lemma 4.1. *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 $V = \{v_1, v_2, \dots, v_n\}$. 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 $C_0 = \{v_1, v_2, \dots, v_{k+2}, v_1\}$. 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 ℓ^{\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.

4.3 Welfare Guarantees with Revealed Preference Rules

Our results from Section 3 give approximate rationality guarantees on individual bidders. We briefly outline this here. By applying the above relaxed revealed preference bidding rules to each bidder, we can now obtain guarantees on the overall social welfare of the entire auction. For example, consider imposing the relaxed GARP bidding rules. Now suppose each bidder uses a minimum, individually rational, ϵ -approximate virtual valuation function

that satisfies the gross substitutes property. It is known that if bidder valuation functions satisfy the gross substitutes property then the simultaneous multi-round auction (SMRA) will converge to a Walrasian equilibrium and maximize social welfare [17,12,16]. Consequently, the output allocation now maximizes virtual welfare to within an additive factor per bidder. One may expect there is some maximum discrepancy (say, in the L^∞ norm) between the true valuation function and *some* virtual approximate virtual valuation. If so, because the implemented virtual valuations are minimum, we can then lower-bound the true social welfare. Similarly, best-case bounds follow using the maximum approximate virtual valuation function.

4.4 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.

Revealed Preference with Budgets. Recall that, in Section 2.2, we have assumed that, in the quasilinear model, bidders have no budgetary constraints. This is not a natural assumption. Harsha et al. [13] 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 $\mathbf{x}_t \succeq \mathbf{y}$ only if $\mathbf{p}_t \cdot \mathbf{y} \leq \lambda \cdot \mathbf{p}_t \cdot \mathbf{x}_t$ where $\lambda < 1$. This is equivalent, in Afriat's original setting, to removing from the graph any arc $(\mathbf{x}_t, \mathbf{x}_s)$ for which $\mathbf{p}_t \cdot \mathbf{x}_s > \lambda \cdot \mathbf{p}_t \cdot \mathbf{x}_t$. Of course, for the application of combinatorial auctions, we assume quasi-linear utilities. Therefore, the appropriate implementation is to remove any arc $(\mathbf{x}_t, \mathbf{x}_s)$ for which

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

How, though, can we implement this rule as $v()$ is unknown? We can simply apply the techniques of Section 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 analog 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 \mathbf{x}_s , where $s < t$. This may be important strategically. To see this, suppose the bid \mathbf{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 $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_t\}$. If \mathbf{x}_s lies on all such negative cycles then \mathbf{x}_t becomes a valid bid after the withdrawal of \mathbf{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 $k + 1$.

References

1. S. Afriat, “The construction of a utility function from expenditure data”, *International Economic Review*, **8**, pp67–77, 1967.
2. G. Akerlof and J. Yellen, “Can small deviations from rationality make significant differences to economic equilibria?”, *Amer. Econ. Rev.*, **75(4)**, pp708–720, 1985.
3. L. Ausubel and O. Baranov, “Market design and the evolution of the combinatorial clock auction”, *Amer. Econ. Rev.*, **104(5)**, pp446–451, 2014.
4. 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.
5. D. Brown and F. Echenique, “Supermodularity and preferences”, *Journal of Economic Theory*, **144(3)**, pp1004–1014, 2009.
6. P. Cramton, “Spectrum auction design”, *Review of Industrial Organization*, **42(2)**, pp161–190, 2013.

7. F. Echenique, D. Golovin and A. Wierman, “A revealed preference approach to computational complexity in economics”, *Proceedings of EC*, pp101–110, 2011.
8. B. Edelman, M. Ostrovsky and Schwarz, “Internet advertising and the generalized second-price auction: selling billions of dollars worth of keywords”, *Amer. Econ. Rev.*, **97(1)**, pp242–259, 2007.
9. A. Fostel, H. Scarf and M. Todd, “Two new proofs of Afriat’s theorem”, *Economic Theory*, **24**, pp211–219, 2004.
10. G. Gigerenzer and R. Selten (eds), *Bounded Rationality: the Adaptive Toolbox*, MIT Press, 2001.
11. J. Gross, “Testing data for consistency with revealed preference”, *The Review of Economics and Statistics*, **77(4)**, pp701–710, 1995.
12. F. Gul and E. Stacchetti, “Walrasian equilibrium with gross substitutes”, *Journal of Economic Theory*, **87**, pp95–124, 1999.
13. 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.
14. H. Houthakker, “Revealed Preference and the Utility Function”, *Economica, New Series*, **17(66)**, pp159–17, 1950.
15. R. Karp, “A characterization of the minimum cycle mean in a digraph”, *Discrete Mathematics*, **23(3)**, p.309–311, 1978.
16. A. Kelso and P. Crawford, “Job matching, coalition formation, and gross substitutes”, *Econometrica*, **50(6)**, pp1483–1504, 1982.
17. P. Milgrom, “Putting auction theory to work: the simultaneous ascending auction”, *Journal of Political Economy*, **108**, pp245–272, 2000.
18. P. Samuelson, “A note on the pure theory of consumer’s behavior”, *Economica*, **5(17)**, pp61–71, 1938.
19. P. Samuelson, “Consumption theory in terms of revealed preference”, *Economica*, **15(60)**, pp243–253, 1948.
20. H. Varian, “Revealed preference”, in Szenberg, Ramrattand and Gottesman (eds.), *Samulesonian Economics and the 21st Century*, pp99–115, O.U.P., 2005.
21. H. Varian, “The nonparametric approach to demand analysis”, *Econometrica*, **50(4)**, pp945–973, 1982.
22. H. Varian, “Goodness-of-fit in optimizing models”, *Journal of Econometrics*, **46**, pp125–140, 1990.
23. H. Varian, “Position auctions”, *Int. J. Ind. Organ.*, **25(6)**, pp1163–1178, 2007.
24. H. Varian, “Revealed preference and its applications”, working paper, 2011.
25. R. Vohra, *Mechanism Design: A Linear Programming Approach*, C.U.P., 2011.