

Stepwise Randomized Combinatorial Auctions Achieve Revenue Monotonicity

Baharak Rastegari*

Anne Condon†

Kevin Leyton-Brown‡

Abstract

In combinatorial auctions that use VCG, a seller can sometimes increase revenue by dropping bidders (see e.g. [5]). In our previous work [26], we showed that such failures of “revenue monotonicity” occur under an extremely broad range of deterministic strategyproof combinatorial auction mechanisms, even when bidders have “known single-minded” valuations. In this work we consider the question of whether revenue monotonic, strategyproof mechanisms for such bidders can be found in the broader class of randomized mechanisms. We demonstrate that—surprisingly—such mechanisms do exist, show how they can be constructed, and consider algorithmic techniques for implementing them in polynomial time.

More formally, we characterize a class of randomized mechanisms defined for known single-minded bidders that are strategyproof and revenue monotonic, and furthermore satisfy some other desirable properties, namely participation, consumer sovereignty and maximality, representing the mechanism as a solution to a quadratically constrained linear program (QCLP). We prove that the QCLP is always feasible (i.e., for all bidder valuations) and give its solution analytically. Furthermore, we give an algorithm for running such a mechanism in time polynomial in the number of bidders and goods; this is interesting because constructing an instance of such mechanisms from our QCLP formulation in a naive way can require exponential time.

1 Introduction

In combinatorial auctions, multiple goods are sold simultaneously and bidders are allowed to place bids on bundles, rather than just on individual goods. These auctions have been widely studied in the last decade, with the ultimate goal of better allocating scarce resources among bidders who value them non-additively (see e.g. [9]). When designing a combinatorial auction mechanism, one may desire that it satisfy various different properties. One important property is that it be a dominant strategy for selfish bidders to truthfully reveal their private information to the mechanism.

Quite a lot of work in the literature is concerned with the design of strategyproof mechanisms for combinatorial auctions [2, 3, 6, 7, 17, 19, 25, 31, 32]. Another important class of properties concerns an auction’s revenue. An auction mechanism is called *optimal* if it maximizes the expected revenue. Optimal auctions were originally studied in the context of single-good auctions [16, 23, 28]. More recent work has extended these ideas to design multi-unit or multi-good auctions that offer strong revenue guarantees, usually achieving a constant fraction of the optimal revenue [1, 8, 14, 20, 29].

We are concerned with describing the way an auction’s revenue changes with the number of participating bidders. Intuitively, one might expect that revenue weakly increases as the number of bidders grows, as competition also increases. We say that an auction mechanism is *revenue monotonic* when this intuition is correct: the seller’s revenue is guaranteed to weakly increase as bidders are added. Groves mechanisms in general and VCG in particular have gained substantial attention because they are the only strategyproof mechanisms that guarantee efficient allocations [15]. However, VCG has also received numerous criticisms ([5, 30]). One of these problems is that VCG is not revenue monotonic for bidders (unless bidders’ valuations are restricted; [4]). Following an example due to [5], consider an auction with three bidders and two goods for sale. Suppose that bidder 2 wants both goods for the price of \$2 billion whereas bidder 1 and bidder 3 are willing to pay \$2 billion for the first and the second good respectively. The VCG mechanism awards the goods to bidders 1 and 3 for the price of zero, yielding the seller zero revenue. However, in the absence of either bidder 1 or bidder 3, the revenue of the auction would be \$2 billion. In our previous work [26, 27] we showed that this problem is not restricted to VCG. Instead, we proved that no revenue monotonic mechanism exists in a very broad class of deterministic, strategyproof combinatorial auction mechanisms. We define this setting and class of mechanisms in Section 2, and also state our impossibility result. Here, we note two lines of research that are closely related to our own past work. First, Day and Milgrom [10, 21] used coalitional game theory to investigate revenue monotonicity in the context of efficient mechanisms. Second, Yokoo et al. [31, 32] investigated false-name bidding; however, their proof can also be understood as showing that revenue monotonicity fails in ef-

*Dept. of CS, University of British Columbia. baharak@cs.ubc.ca.

†Dept. of CS, University of British Columbia. condon@cs.ubc.ca.

‡Dept. of CS, University of British Columbia. kevinlb@cs.ubc.ca

efficient strategyproof mechanisms. In contrast, we do not restrict ourselves to efficient mechanisms.

There are many cases ([11, 12, 13, 18]) in which randomized mechanisms are able to achieve desirable properties that cannot be obtained by deterministic mechanisms. In Section 3 we define randomized mechanisms and some desirable properties for our setting. In Section 4 we show that it is possible to circumvent our impossibility result, at least for known single-minded bidders, by proposing a class of combinatorial auction mechanisms that we call “stepwise” randomized combinatorial auction mechanisms. We also show how to construct such revenue-monotonic mechanisms, though this construction can sometimes require exponential time. Finally, in Section 5 we give a polynomial-time algorithm for constructing our mechanism.

2 Deterministic Mechanisms

To prove results about revenue monotonicity, we need to reason about the behaviour of combinatorial auction mechanisms when bidders are added or dropped. We also need to reason about mechanisms whose behavior can depend on bidder preferences—for example, each bidder may have a “single-minded” interest in one particular bundle. For these reasons, we provide a set of general definitions in which the allocation of goods and the payments imposed may depend on which bidders participate and which goods are for sale, as well as on bidders’ declared preferences.

Let $\mathbb{N} = \{1, \dots, n\}$ be the universal set of n bidders—all the potential bidders who exist in the world. Let $N \subseteq \mathbb{N}$ denote the set of bidders participating in a particular auction. Let \mathbb{G} be the finite universe of goods for sale. Let $G \subseteq \mathbb{G}$ denote the set of goods for sale in a particular auction. Let both N and G be common knowledge among all bidders and the auctioneer.

A *valuation function* describes the values that a bidder holds for subsets of the set of goods in G . Let valuation function $v_{G,i}$ for bidder $i \in \mathbb{N}$ map 2^G to the nonnegative reals. For every $G \subseteq \mathbb{G}$ let valuation function $v_{G,i}$ be the projection of $v_{\mathbb{G},i}$ into G . Whenever G is understood, we drop it from the subscript. We assume that bidders have quasilinear utility functions; that is, bidder i ’s utility for bundle a_i is $v_i(a_i) - p_i$, where v_i is her valuation and p_i is any payment she is required to make.

A *valuation profile* is an n -tuple $v = (v_1, \dots, v_n)$, where, for every participating bidder i , v_i is a valuation function. Let \mathbb{V} denote the universal set of all possible valuation profiles. Observe that valuation profiles always have one entry for every potential bidder, regardless of the number of bidders who participate in the auction. We use the symbol \emptyset in such tuples as a placeholder for each non-participating bidder (i.e., each bidder $i \notin N$). When v is an n -dimensional tuple, then $(v_1, \dots, v_{i-1}, \emptyset, v_{i+1}, \dots, v_n)$ is denoted by v_{-i} . Note that if $i \notin N$, then $v = v_{-i}$. Let

$\mathbb{V}_{N,G}$ denote the set of all valuation profiles given a set of participating bidders N and a set of goods for sale G ; that is, the set of all valuation profiles v_G for which $v_i = \emptyset$ if and only if $i \notin N$.

If asked to reveal her valuation, a bidder may not tell the truth. Denote the declared valuation function of a (participating) bidder i as \hat{v}_i . Let \hat{v} be the declared valuation profile. Use the same notation to describe declared valuation profiles as valuation profiles (e.g., all declared valuation profiles are n -tuples), and furthermore write (v_i, \hat{v}_{-i}) to denote $(\hat{v}_1, \dots, \hat{v}_{i-1}, v_i, \hat{v}_{i+1}, \dots, \hat{v}_n)$.

In a particular auction, bidders’ valuation functions may be drawn from some restricted set. For example, we will need to make such an assumption to model known single-minded bidders. Let $V_{N,G} \subseteq \mathbb{V}_{N,G}$ denote a subspace of the universal set of valuation profiles for the set of participating bidders N and the set of goods for sale G . (For example, all valuations consistent with each bidder having a single-minded interest in one known bundle.) Let $\mathcal{V}_{\mathbb{N},\mathbb{G}}$ denote the universal set of valuation profile subspaces, that is $\mathcal{V}_{\mathbb{N},\mathbb{G}} = \{V_{N,G} \mid N \subseteq \mathbb{N}, G \subseteq \mathbb{G}, V_{N,G} \subseteq \mathbb{V}_{N,G}\}$. Let \mathcal{V} denote a set of valuation profile subspaces with at least one member corresponding any $N \subseteq \mathbb{N}$ and $G \subseteq \mathbb{G}$. That is, $\mathcal{V} \subseteq \mathcal{V}_{\mathbb{N},\mathbb{G}}$ and $\exists V_{N,G} \in \mathcal{V}, \forall N \subseteq \mathbb{N}, G \subseteq \mathbb{G}$. (For example, subspaces corresponding to all the possible sets of known bundles for different bidders.) Note that there could be more than one subspace corresponding to a fixed N and a fixed G in \mathcal{V} .

DEFINITION 2.1. (CA MECHANISM) *Let set of valuation profile subspaces \mathcal{V} be given. A deterministic direct Combinatorial Auction (CA) mechanism M (CA mechanism) maps each $V_{N,G} \in \mathcal{V}$, $N \subseteq \mathbb{N}$ and $G \subseteq \mathbb{G}$, to a pair (a, p) where*

- *a , the allocation scheme, maps each $\hat{v} \in V_{N,G}$ to an allocation tuple $a = (a_1(\hat{v}), \dots, a_n(\hat{v}))$ of goods, where $\cup_i a_i(\hat{v}) \subseteq G$, $a_i(\hat{v}) \cap a_j(\hat{v}) = \emptyset$ if $i \neq j$, and $a_i(\hat{v}) = \emptyset$ if $\hat{v}_i = \emptyset$.*
- *p , the payment scheme, maps each $\hat{v} \in V_{N,G}$ to a payment tuple $p = (p_1(\hat{v}), \dots, p_n(\hat{v}))$, where $p_i(\hat{v})$ is the payment from bidder i to the auctioneer such that $p_i(\hat{v}) = 0$ if $\hat{v}_i = \emptyset$.*

We refer to a_i and p_i as bidder i ’s allocation and payment functions respectively. Whenever \hat{v} can be understood from the context, we refer to $a_i(\hat{v})$ and $p_i(\hat{v})$ by a_i and p_i , respectively. If $\hat{v}_i(a_i) > 0$, we say that bidder i “wins”. We denote by \mathbb{A}_G the set of all possible partitions of G into n partitions; i.e. the set of all possible ways of distributing goods among participating bidders. For any given allocation $a \in \mathbb{A}_G$, we denote by a_i the set of goods that are allocated to bidder i under a .

Mechanisms that give rise to dominant strategies are especially desirable, as bidders are spared having to reason about each others’ behavior. A direct CA mechanism is

said to be *truthful* if in equilibrium bidders declare their true valuations to the mechanism. A direct CA mechanism is said to be *strategyproof* (or *dominant strategy truthful*) if every bidder has the dominant strategy of revealing her true preferences.

The revenue of an auction is the sum of payments made to the auctioneer. Informally, an auction mechanism is revenue monotonic if the auctioneer could never increase revenue by dropping a bidder.

DEFINITION 2.2. (REVENUE MONOTONICITY) *A truthful CA mechanism M is revenue monotonic if and only if for all $N \subseteq \mathbb{N}$, $G \subseteq \mathbb{G}$, $V_{N,G} \in \mathcal{V}$, $v \in V_{N,G}$ and for all bidders j ,*

$$\sum_{i \in \mathbb{N}} p_i(v) \geq \sum_{i \in \mathbb{N} \setminus \{j\}} p_i(v_{-j}).$$

It is natural, and commonly assumed, that a bidder should make no payment to the mechanism unless she wins. We call this assumption *participation*. A mechanism is *weakly maximal* with respect to a bidder i if, whenever i values any good g sufficiently, the mechanism does not withhold that good or give it away to a bidder who does not value it. (For formal definitions see [27].)

Our main results refer to a restricted class of valuation spaces: known single-minded bidders. Our definition of this class follows Mu’alem and Nisan [22] and Nisan [24]. Informally, a participating bidder i is single-minded if she only values bundles that contain a particular bundle b_i , and she values all these bundles equally. Thus, i ’s valuation function v_i maps supersets of b_i to some positive value v_i and maps all other bundles to 0. A mechanism is defined for known single-minded bidders if all bidders are single-minded and furthermore, the mechanism “knows” the bundle b_i that is valued by each participating bidder i . Thus, bidder i can lie about her value (declare a $\hat{v}_i \neq v_i$), but does not even have to declare her bundle of interest b_i , and hence cannot lie about it.

More formally, let $b = (b_1, b_2, \dots, b_n) \in (2^G)^n$. Fix $N \subseteq \mathbb{N}$ and $G \subseteq \mathbb{G}$. If i is a participating bidder, let $V_{N,G,i}^{(b)}$ be the set of all such functions, taken over all possible choices of v_i , and otherwise let $V_{N,G,i}^{(b)} = \emptyset$. Let $V_{N,G}^{(b)} = V_{N,G,1}^{(b)} \times \dots \times V_{N,G,n}^{(b)}$. Thus, $V_{N,G}^{(b)}$ is simply the space of valuation profiles in which participating bidders are all single-minded, with the bundle valued by participating bidder i being b_i .

Let $\mathcal{V}^{(ksm)}$ denote the set of valuation profile subspaces for known single-minded bidders; that is $\mathcal{V}^{(ksm)} = \{V_{N,G}^{(b)} \mid N \subseteq \mathbb{N}, G \subseteq \mathbb{G}, b \in (2^G)^n\}$. We say that a mechanism is *defined for known single-minded bidders* if its set of valuation profile subspaces is $\mathcal{V}^{(ksm)}$. From the definition of mechanism (Definition 2.1), it follows that the allocation and payment functions depend on the set $V_{N,G}^{(b)} \in \mathcal{V}^{(ksm)}$ from which bidders’ valuation profiles are drawn. Informally, b is *known*, since the allocation and payments depend on b .

The valuation of a known single-minded bidder can be characterized by the single parameter v_i , representing i ’s valuation for any superset of bundle b_i . Thus in this case we use v to denote single-minded bidders’ valuation profile, \hat{v}_i to denote the declared valuation of a participating bidder i , and \hat{v} to denote a tuple consisting of declared valuations for each participating bidder and \emptyset symbols for each non-participating bidder.

Roughly speaking, a mechanism defined for known single-minded bidders satisfies *consumer sovereignty* if by bidding high enough, any bidder can win the bundle she values—more formally, given any bidder i and the declared values of the other bidders, \hat{v}_{-i} , there exists some finite amount $k_i \in \mathbb{R}$, $k_i > 0$, such that if i reports $\hat{v}_i = k_i$ then i is allocated at least b_i .

In our past work [26, 27], we proved an impossibility result: there is no deterministic combinatorial auction mechanism that satisfies our desirable properties. To obtain as strong a result as possible, we proved that the result is true even when the bidders are known single-minded.

THEOREM 2.1. *Let $|\mathbb{G}| \geq 2$ and $|\mathbb{N}| \geq 3$. Let M be a CA mechanism defined for known single-minded bidders that offers dominant strategies to bidders and satisfies participation, consumer sovereignty, and weak maximality with respect to at least two bidders. Then M is not revenue monotonic.*

3 Randomized Mechanisms

In this work, we study the consequences of relaxing the assumption that mechanisms are deterministic. More specifically, we ask whether there are revenue monotonic *randomized* mechanisms that satisfy our desired properties. (As we will see, moving to the randomized case will also require reinterpretations of these properties, particularly of consumer sovereignty.) As the following definition states, a randomized CA mechanism produces a distribution over allocations and payments.

DEFINITION 3.1. (RANDOMIZED CA MECHANISM (RCAM)) *Let \mathcal{V} be given. A randomized direct Combinatorial Auction mechanism M (RCAM) maps each $V_{N,G} \in \mathcal{V}$, $N \subseteq \mathbb{N}$ and $G \subseteq \mathbb{G}$, to a distribution over pairs (a, p) where a and p are defined exactly as in Definition 2.1.*

Given $V_{N,G} \in \mathcal{V}$, let $\pi_a(\hat{v})$ denote the probability that allocation $a \in \mathbb{A}_G$ will be chosen given declared values \hat{v} . Let $p_i(\hat{v})$ denote the expected payment of i .

A randomized CA mechanism is *strategyproof in expectation* if and only if truth-telling is a dominant strategy for all bidders in the game induced by expectation.

Randomized mechanisms can be defined for known single-minded bidders in a manner analogous to that used for deterministic mechanisms above. In what follows, we

concern ourselves only with randomized mechanisms for known single-minded bidders.

For a randomized CA mechanism that is defined for known single-minded bidders, let $w_i(\hat{v})$ denote the probability that bidder i wins—that is, i is allocated a bundle that includes b_i , given $V_{N,G}^{(b)}$. Note that the $\pi_a(\hat{v})$'s fully define $w_i(\hat{v})$'s. Formally,

$$(3.1) \quad w_i(\hat{v}) = \sum_{\forall a \in \mathbb{A}_G, a_i \supseteq b_i} \pi_a(\hat{v}).$$

The following theorem characterizes the class of strategyproof randomized mechanisms defined over known single-minded bidders (indeed, for any single parameter domain).

THEOREM 3.1. (SEE E.G. [24]) *A randomized mechanism defined over known single-minded bidders is strategyproof in expectation, and satisfies participation, iff for all $V_{N,G}^{(b)}$ and every bidder $i \in N$ and every fixed \hat{v}_{-i} we have that*

1. *the function $w_i(\hat{v}_i, \hat{v}_{-i})$ is monotonically non decreasing in \hat{v}_i .*
2. *$p_i(\hat{v}_i, \hat{v}_{-i}) = \hat{v}_i \cdot w_i(\hat{v}_i, \hat{v}_{-i}) - \int_{t=0}^{t=\hat{v}_i} w_i(t, \hat{v}_{-i}) dt$.*

COROLLARY 3.1. (IMMEDIATE FROM THEOREM 3.1) *A strategyproof mechanism satisfies participation if and only if it is characterized by a set of feasible allocation distributions $\pi_a(\hat{v})$'s that induce monotonic winning probability functions $w_i(\hat{v})$'s and p_i 's are defined as in Theorem 3.1.*

Now we generalize the properties we defined for deterministic mechanisms to the randomized setting.

DEFINITION 3.2. (REVENUE MONOTONIC FOR RCAM'S) *A truthful randomized CA mechanism is revenue monotonic if dropping a bidder never increases the mechanism's expected revenue.*

DEFINITION 3.3. (PARTICIPATION FOR RCAM'S) *A truthful randomized CA mechanism satisfies participation iff for all $N \subseteq \mathbb{N}, G \subseteq \mathbb{G}, V_{N,G} \in \mathcal{V}$, and $v \in V_{N,G}$, $p_i(v) = 0$ for any bidder i for whom $w_i(v) = 0$.*

Next we define maximality for randomized CA mechanisms. Our definition here is stronger (more general) than the weak maximality definition we provided for deterministic mechanisms. We chose to use weak maximality in [26, 27] because it strengthened our impossibility result stated there. For randomized mechanisms, we will prove a positive result, namely the existence of revenue monotonic randomized mechanisms with several desired properties. To make our positive result as strong as possible, we use a more general notion of maximality here. Informally, a mechanism is maximal with respect to a bidder i if, whenever i values any *subset* of goods s (rather than a single good g) sufficiently, the mechanism does not withhold that bundle or give the goods in the bundle away to bidders who do not value them.

DEFINITION 3.4. (MAXIMALITY FOR RCAM'S) *A truthful randomized CA mechanism M is maximal with respect to bidder i iff $\forall N \subseteq \mathbb{N}$ and $\forall G \subseteq \mathbb{G}$ there exists a set of nonnegative finite constants $\{\alpha_{N,G,i,s} \mid s \subseteq G\}$ such that the following holds. For all $i \in N$, $V_{N,G} \in \mathcal{V}$, and $v \in V_{N,G}$, for any allocation a that has a positive support under M —that is, a is chosen by M with probability above zero—either:*

1. *$v_i(a_i) > 0$; or*
2. *for any allocation a' with $v_i(a'_i) > \alpha_{N,G,i,a'_i}$ and $a'_j = a_j \setminus a'_i$ for all $j \neq i$, it must be the case that for some j , $v_j(a'_j) < v_j(a_j)$.*

An allocation a is maximal if it satisfies either (1) or (2) for all bidders i . A randomized CA mechanism M satisfies maximality if any allocation with positive support under M is maximal.

It is somewhat harder to decide how to extend our consumer sovereignty definition to randomized mechanisms for known single-minded bidders. We first consider two possible extensions to the definition for deterministic mechanisms, which can be seen as opposite extremes. First, we could define consumer sovereignty (I) as requiring that, fixing bids of the others, any bidder is able to win any desired bundle with probability one if she bids high enough. Unfortunately, in this case we recover our previous impossibility result.

THEOREM 3.2. (INFORMAL) *Let $|\mathbb{G}| \geq 2$ and $|\mathbb{N}| \geq 3$. Let M be a randomized CA mechanism defined for known single-minded bidders that offers dominant strategies to bidders and satisfies participation, consumer sovereignty (I), and weak maximality with respect to at least two bidders. Then M is not revenue monotonic.*

On the other hand, we could define consumer sovereignty (II) as requiring that any bidder be able to win any desired bundle with *some* probability above zero if she bids high enough. This leads to a different problem. Consider a mechanism M with $\alpha_{N,G,i,s} = 0$ that chooses a maximal allocation uniformly at random, and charges nothing. Note that each bidder wins her desired bundle in at least one maximal allocation. Therefore, it is easy to verify that M is strategyproof and satisfies participation, consumer sovereignty, and maximality with respect to all bidders. It also is revenue monotonic since it never collects any money.

The above arguments suggest that we ought to seek an intermediate definition for consumer sovereignty. We thus present the following definition, which roughly requires that, given the valuations of the other bidders, a bidder who starts bidding at 0 and then raises her bid can increase her probability of winning by at least δ at least γ times.

DEFINITION 3.5. ((γ -STEP, δ) CONSUMER SOVEREIGNTY) *A randomized CA mechanism defined for known*

single-minded bidders satisfies $(\gamma\text{-step}, \delta)$ consumer sovereignty, $\gamma \geq 0$ and $\delta > 0$, iff for any fixed tuple of bundles $b = (b_1, \dots, b_n)$ and for some constants $0 = c_{i,0} < c_{i,1} < \dots < c_{i,\gamma} < c_{i,\gamma+1} = \infty$, $\forall i \in \mathbb{N}$, the following holds. For all $N \subseteq \mathbb{N}$, $G \subseteq \mathbb{G}$, bidder $i \in N$, $\hat{v}_{-i} \in \left(V_{N,G}^{(b)}\right)_{-i}$, and $j < \gamma$, we have that: the winning probabilities, w_i 's, are monotonic and furthermore either $w_i(c_{i,s_i+1}, \hat{v}_{-i}) \geq w_i(c_{i,s_i}, \hat{v}_{-i}) + \delta$ or $w_i(c_{i,s_i+1}, \hat{v}_{-i}) = 1$.

It is easy to see that if a mechanism satisfies $(\gamma\text{-step}, \delta)$ consumer sovereignty for some $\gamma = k$, it then also satisfies $(\gamma\text{-step}, \delta)$ consumer sovereignty for any γ for which $0 \leq \gamma \leq k$. Observe that the constants c_{i,s_i} are independent of all bidders' declared valuations; in a sense, they can be seen as "bidder-specific, leveled reserve prices." Thus, while we do not assume that the mechanism designer knows anything about the valuation distribution(s), if such information is available, it can be useful for setting these constants.

We now propose a simple and useful class of randomized mechanisms. These define the probability that any given bidder wins as a stepwise function of her bid amount, with a finite number of steps.

DEFINITION 3.6. (STEPWISE RANDOMIZED MECHANISM)
A randomized CA mechanism defined for known single-minded bidders is a stepwise mechanism if for some $k > 0$ and some constants $0 = c_{i,0} < c_{i,1} < \dots < c_{i,k} < c_{i,k+1} = \infty$, $\forall i \in \mathbb{N}$, the following holds. For all fixed tuples of bundles $b = (b_1, \dots, b_n)$, for all $N \subseteq \mathbb{N}$, $G \subseteq \mathbb{G}$, for all bidders $i \in N$ and valuation profiles $\hat{v}_{-i} \in \left(V_{N,G}^{(b)}\right)_{-i}$, and for all $c_{i,s_i} \leq \hat{v}_i < c_{i,s_i+1}$, it is the case that $w_{N,\ell}(\hat{v}_i, \hat{v}_{-i}) = w_{N,\ell}(c_{i,s_i}, \hat{v}_{-i})$, for all $\ell \in N$.

We call the mechanism a γ -step randomized mechanism if it satisfies the above for $k = \gamma$.

A γ -step randomized mechanism can be interpreted as a mechanism that for each bidder i , cares only about specific declared values, $c_{i,0}, c_{i,1}, \dots, c_{i,\gamma}$ and treats any declared value of i between c_{i,s_i} and c_{i,s_i+1} the same as c_{i,s_i} . In fact, one can easily verify that $w_i(\hat{v}) = w_i(c_{1,s_1}, \dots, c_{n,s_n})$, for all \hat{v} where $c_{i,s_i} \leq \hat{v}_i < c_{i,s_i+1}$ for all participating bidders i . If a γ -step randomized mechanism additionally has the following monotonicity property that either (1) $w_i(c_{i,s_i}, \hat{v}_{-i}) + \delta \leq w_i(c_{i,s_i+1}, \hat{v}_{-i})$, or (2) $w_i(c_{i,s_i+1}, \hat{v}_{-i}) = 1$, then the mechanism satisfies $(\gamma\text{-step}, \delta)$ consumer sovereignty.

Figure 1 shows a sample w_i for a 6-step stepwise randomized mechanism, given fixed bids by the other bidders. Observe that, by Theorem 3.1, if the mechanism is to satisfy strategyproofness and participation, our choice of w_i must imply a specific choice of p_i . Here, if the bidder declares \hat{v}_i , she must pay an amount equal to the area of the shaded region.




Figure 1: i 's probability of winning as a function of her bid amount, given fixed bids by the other agents.

4 A Revenue Monotonic Mechanism

In this section, we construct a γ -step randomized mechanism, which we dub M_γ , that is strategyproof and revenue monotonic and satisfies participation, maximality and $(\gamma\text{-step}, \delta)$ consumer sovereignty, for any given γ and for some $\delta > 0$. We construct M_γ such that when a bidder i increases her bid one step, her probability of winning increases by δ unless she wins in all maximal allocations, in which case her probability of winning is equal to 1. We first give a nonlinear feasibility program F and show that its solutions correspond to mechanisms that satisfy all our desired properties. We then construct a quadratically constrained linear program (QLCP) P , and prove that all of its solutions that satisfy one additional constraint also solve F . Finally, we constructively prove that such solutions of P always exist.

Given $V_{N,G}^{(b)}$, for all $N \subseteq \mathbb{N}$ let \mathbb{M}_N be the set of all maximal allocations with respect to maximality parameters set to zero—that is, $\alpha_{N,G,i,s} = 0, \forall i \in N, s \in G$. Let \mathcal{M}_N be a set of maximal allocations—that is $\mathcal{M}_N \subseteq \mathbb{M}_N$ —such that each bidder is either allocated her desired bundle or nothing and such that each bidder wins in at least one allocation $a \in \mathcal{M}_N$. Let $\hat{\theta}$ denote the tuple of declared valuations in which all participating bidders bid 0.

LEMMA 4.1. For all $V_{N,G}^{(b)}$, all $N' \subseteq N \subseteq \mathbb{N}$, and all bidders $i \in N'$, if i belongs to all allocations $a \in \mathcal{M}_N$ then i belongs to all allocations $a' \in \mathcal{M}_{N'}$.

Proof. Since i belongs to all $a \in \mathcal{M}_N$, then it must be the case that b_i does not overlap with any other bidder's desired bundle; that is $b_i \cap b_j = \emptyset, \forall j \in N, j \neq i$. Therefore, since $N' \subseteq N$, it is also true that $b_i \cap b_j = \emptyset, \forall j \in N', j \neq i$. Thus, i has to belong to all maximal allocations under N' and therefore to all allocations $a' \in \mathcal{M}_{N'}$. \square

4.1 Feasibility program Observe that a mechanism is a mapping from declared valuations to allocation probabilities π_a 's and payments p_i 's. Here we express such a mapping as a solution to a set of feasibility programs (albeit ones with

$$\begin{aligned}
(F.a_1) \quad & 0 \leq \pi_{N,a}(\hat{v}) \leq 1 && \forall N, \hat{v}, a \in \mathbb{A}_G \\
(F.a_2) \quad & \sum_{a \in \mathbb{A}_G} \pi_{N,a}(\hat{v}) = 1 && \forall N, \hat{v} \\
(F.w) \quad & w_{N,i}(\hat{v}) = \sum_{a \in \mathbb{A}_G, a_i \supseteq b_i} \pi_{N,a}(\hat{v}) && \forall N, i, \hat{v} \\
(F.step) \quad & w_{N,l}(\hat{v}_i, \hat{v}_{-i}) = w_{N,l}(c_{i,s_i}, \hat{v}_{-i}) && \forall N, i, l, s_i, \hat{v} | c_{i,s_i} \leq \hat{v}_i < c_{i,s_i+1} \\
(F.mon) \quad & w_{N,i}(\hat{v}) \geq w_{N,i}(\hat{v}') && \forall N, i, \hat{v}, \hat{v}' | \hat{v}_i \geq \hat{v}'_i \text{ and } \hat{v}_{-i} = \hat{v}'_{-i} \\
(F.sp) \quad & p_{N,i}(\hat{v}_i, \hat{v}_{-i}) = \hat{v}_i \cdot w_{N,i}(\hat{v}_i, \hat{v}_{-i}) - \int_{t=0}^{t=\hat{v}_i} w_{N,i}(t, \hat{v}_{-i}) dt && \forall N, i, \hat{v} \\
(F.max) \quad & \pi_{N,a}(\hat{v}) = 0 && \forall N, \hat{v}, a \notin \mathbb{M}_N \\
(F.rm) \quad & \sum_i p_{N,i}(\hat{v}) \geq \sum_{i \neq l} p_{N \setminus \{l\}, i}(\hat{v}_{-l}) && \forall N, l, \hat{v} \\
(F.cs) \quad & w_{N,i}(c_{i,s_i+1}, \hat{v}_{-i}) \geq w_{N,i}(c_{i,s_i}, \hat{v}_{-i}) + \delta \text{ or } w_{N,i}(c_{i,s_i+1}, \hat{v}_{-i}) = 1 && \forall N, i, s_i, \hat{v}_{-i} \in \left(V_{N,G}^{(b)} \right)_{-i} \\
(F.\delta) \quad & \delta > 0 &&
\end{aligned}$$

Figure 2: Nonlinear feasibility program $F(V_{\mathbb{N},G}^{(b)}, G)$. Constants are \hat{v} 's and c_{i,s_i} 's. Variables are $\pi_{N,a}$'s, $w_{N,i}$'s, $p_{N,i}$'s and δ . We adopt the conventions that N indexes subsets of \mathbb{N} , i and l index elements of N , s_i indexes elements of $\{0, \dots, \gamma\}$, and \hat{v} indexes elements of $V_{N,G}^{(b)}$. Observe that because this last set is (uncountably) infinite, the feasibility program involves an infinite number of both variables and constraints.

some nonlinear constraints, and an uncountably infinite number of both variables and constraints). Recall that any CA mechanism defined for known single-minded bidders is able to condition its behavior on G, N , and $V_{N,G}^{(b)}$ (see Definition 2.1). Because the mechanism is free to behave differently for every G (available set of goods) and $V_{N,G}^{(b)}$ (set of known bundles of interest for the bidders), we write a separate feasibility program for each possible joint assignment to these variables. Our feasibility program, denoted F and given in Figure 2, is thus parameterized by $V_{\mathbb{N},G}^{(b)}$ and G . Note that we have introduced the assumption that the mechanism knows $V_{\mathbb{N},G}^{(b)}$ rather than $V_{N,G}^{(b)}$ (i.e., it knows the bundles of non-participating bidders). This assumption will make no difference in what follows, but dramatically simplifies notation.

LEMMA 4.2. *Any solution to $F(V_{\mathbb{N},G}^{(b)}, G)$ for all $V_{\mathbb{N},G}^{(b)}$ and $G \subseteq \mathbb{G}$, corresponds to a γ -step randomized mechanism that satisfies strategyproofness, participation, maximality, (γ -step, δ) consumer sovereignty, and revenue monotonicity.*

Proof. We must ensure that a solution to the $F(V_{\mathbb{N},G}^{(b)}, G)$'s induces a valid mechanism. First, it is necessary to ensure that $\pi_{N,a}$'s correspond to probabilities. This is achieved by Constraints (F.a₁) and (F.a₂). Second, Constraint (F.w) ensures that these allocation probabilities fully define winning probabilities, as required by Equation (3.1). Third, Constraint (F.step) ensures that our mechanism is stepwise randomized.

Now we must show that the mechanism satisfies our five desired properties. First, Constraints (F.mon) and (F.sp) together entail both strategyproofness and participation (by

Theorem 3.1). Second, Constraint (F.max) entails maximality. Third, Constraint (F.rm) entails revenue monotonicity. Finally, Constraints (F.mon), (F.cs) and (F. δ) together ensure that the mechanism satisfies (γ -step, δ) consumer sovereignty for a given γ and some $\delta > 0$. \square

4.2 Quadratically constrained linear program Consider quadratically constrained linear program (QLCP) $P(V_{\mathbb{N},G}^{(b)}, G)$ in Figure 3. We will prove that if $P(V_{\mathbb{N},G}^{(b)}, G)$ can be solved for all $V_{\mathbb{N},G}^{(b)}$ and $G \subseteq \mathbb{G}$ with $\delta > 0$ then we can construct solutions for the $F(V_{\mathbb{N},G}^{(b)}, G)$'s and construct our desired mechanism, M_γ . Recall that F is parameterized by an infinite size valuation space, $V_{\mathbb{N},G}^{(b)}$, and thus has an infinite number of variables and constraints. The main idea in this section is that we can move from an infinite-sized F to a finite-sized QLCP P by working with a finite sized valuation space, $\mathbb{V}_{N,G}^{(b)}$. Specifically, for each bidder i we only need to consider the finite set of possible declared values $c_{i,0}, \dots, c_{i,s_\gamma}$. Formally, $\mathbb{V}_{N,G}^{(b)} = \{v | v \in V_{N,G}^{(b)}, \forall i \in N, \exists s_i \in \{0, \dots, \gamma\} : v_i = c_{i,s_i}\}$. To show that any solution of P corresponds to a solution of F we provide a mapping from $\mathbb{V}_{N,G}^{(b)}$ to $V_{N,G}^{(b)}$.

To provide intuition for our proof, we state $P(V_{\mathbb{N},G}^{(b)}, G)$ for a simple example and show how we can find a solution to it that sets $\delta > 0$. Consider the bidder-bundle setting described in the introduction, which we used to demonstrate that VCG is not revenue monotonic. That is, let $\mathbb{G} = \{g_1, g_2\}$ and $\mathbb{N} = \{1, 2, 3\}$; bidders 1, 2 and 3 are known single-minded, where the bundles valued by bidders 1, 2, and 3

maximize δ subject to:

$$\begin{aligned}
(P.\pi_1) \quad & \pi_{N,a}(\hat{\mathcal{V}}) = 0 & \forall N, \hat{\mathcal{V}}, a \in \mathbb{A}_G \setminus \mathcal{M}_N \\
(P.\pi_2) \quad & 0 \leq \pi_{N,a}(\hat{\mathcal{V}}) \leq 1 & \forall N, \hat{\mathcal{V}}, a \in \mathcal{M}_N \\
(P.\pi_3) \quad & \sum_{a \in \mathcal{M}_N} \pi_{N,a}(\hat{\mathcal{V}}) = 1 & \forall N, \hat{\mathcal{V}} \\
(P.\pi_4) \quad & \pi_{N,a}(\hat{\mathcal{V}}) = \pi_{N,a}(\hat{0}) + \sum_i (q_{N,a,i} \cdot \delta \cdot s_i) & \forall N, \hat{\mathcal{V}} | \hat{\mathcal{V}}_i = c_{i,s_i} \\
(P.q_1) \quad & q_{N,a,i} = 0 & \forall N, i, a \in \mathcal{M}_N | \forall a' \in \mathcal{M}_N, a'_i = b_i \\
(P.q_2) \quad & 0 \leq q_{N,a,i} \leq 1 & \forall N, i, a \in \mathcal{M}_N | a_i = b_i \\
(P.q_3) \quad & -1 \leq q_{N,a,i} \leq 0 & \forall N, i, a \in \mathcal{M}_N | a_i = \emptyset \\
(P.q_4) \quad & \sum_{a \in \mathcal{M}_N : a_i = b_i} q_{N,a,i} = 1 & \forall N, i | \exists a' \in \mathcal{M}_N \text{ and } a'_i = \emptyset \\
(P.q_5) \quad & \sum_{a \in \mathcal{M}_N : a_i = \emptyset} q_{N,a,i} = -1 & \forall N, i | \exists a' \in \mathcal{M}_N \text{ and } a'_i = \emptyset
\end{aligned}$$

Figure 3: Quadratically constrained linear program $P(V_{N,G}^{(b)}, G)$. Variables are $\pi_{N,a}$'s, $q_{a,i}$'s and δ . We adopt the conventions that N indexes subsets of \mathbb{N} , i indexes elements of N , s_i indexes elements of $\{0, \dots, \gamma\}$, and $\hat{\mathcal{V}}$ indexes elements of $\mathbb{V}_{N,G}^{(b)} = \{\mathcal{V} | \mathcal{V} \in V_{N,G}^{(b)}, \forall i \in N, \exists s_i \in \{0, \dots, \gamma\} : \mathcal{V}_i = c_{i,s_i}\}$.

are $b_1 = \{g_1\}$, $b_2 = \{g_1, g_2\}$ and $b_3 = \{g_2\}$ respectively. Let $S_{3,2}$ denote this three-bidder, two-good setting. We state the constraints and explain the solution for the case when all bidders are present and all goods are for sale. That is, let $G = \mathbb{G}$ and $N = \mathbb{N}$. One can easily follow the same approach for other choices of G and N , many of which are trivial.

It is easy to verify that there is exactly one choice for \mathcal{M}_N : we have to either award bidder 2 her desired bundle, or award bidders 1 and 3 each their desired bundle. Therefore $\mathcal{M}_N = \{(\emptyset, \{g_1, g_2\}, \emptyset), (\{g_1\}, \emptyset, \{g_2\})\}$. Let $a_2 = (\emptyset, \{g_1, g_2\}, \emptyset)$ and $a_{1,3} = (\{g_1\}, \emptyset, \{g_2\})$. Thus, for all $\hat{\mathcal{V}} \in \mathbb{V}_{N,G}^{(b)}$: (i) $\pi_{N,a}(\hat{\mathcal{V}}) = 0$, for all $a \in \mathbb{A}_G$ such that $a \neq a_2, a_{1,3}$ constitute $(P.\pi_1)$, (ii) $0 \leq \pi_{N,a}(\hat{\mathcal{V}}) \leq 1$ if $a = a_2$ or $a = a_{1,3}$ constitute $(P.\pi_2)$, and (iii) $\pi_{N,a_2}(\hat{\mathcal{V}}) = 1 - \pi_{N,a_{1,3}}(\hat{\mathcal{V}})$ constitute $(P.\pi_3)$.

As each bidder i belongs to exactly one allocation $a \in \mathcal{M}_N$, Constraints $(P.q_1)$ – $(P.q_5)$ can be expressed as $q_{N,a,i} = 0$, for all $a \neq a_2, a_{1,3}$ and all $i \in N$, and $q_{N,a_{1,3},1} = q_{N,a_{1,3},3} = 1$, $q_{N,a_2,2} = 1$, $q_{N,a_{1,3},2} = -1$, and $q_{N,a_2,1} = q_{N,a_2,3} = -1$.

Intuitively, $q_{N,a,i} \cdot \delta$ denotes the change to $\pi_{N,a}$ when bidder i increases her bid by one step. We constrain the $q_{N,a,i}$'s in $(P.\pi_4)$ such that when i increases her bid by one step—from c_{i,s_i} to $c_{i,s_{i+1}}$, the probability that $a \in \mathcal{M}_N$ will be chosen weakly increases if i belongs to a_i , and weakly decreases otherwise.

One can illustrate constraints in $(P.\pi_4)$ by the following graph representation. Let GR_N be a graph of $(\gamma + 1)^{|N|}$ nodes, each corresponding to a different potential declared valuation profile of bidders in $\mathbb{V}_{N,G}^{(b)}$. Let there be a directed edge between each pair of nodes that differ in only one of the bidders' declarations, and in which this difference is an increase of exactly one step (i.e., from c_{i,s_i} to $c_{i,s_{i+1}}$). In

other words, we can move from one node to another by increasing one bidder's bid by one step. If an edge indicates an increase in bidder i 's declared value, we say the edge is of type e_i . Now, assign $|\mathcal{M}_N|$ labels to each edge, one for each allocation in \mathcal{M}_N . Allocation a 's label on an edge of type e_i denotes the change to $\pi_{N,a}$ by moving along an edge of type e_i (which increases bidder i 's bid amount by one step) and is equal to $q_{N,a,i} \cdot \delta$. Define the cost of a path as the absolute value of the sum of the labels of the edges in the path.

LEMMA 4.3. *In each GR_N and for all allocations $a \in \mathcal{M}_N$, all paths between any two given nodes s and t have the same cost: $|\pi_{N,a}(t) - \pi_{N,a}(s)|$.*

Proof. For all $i \in N$, the number of edges of type e_i is the same in all paths between s and t . Since all edges of type e_i have the same label corresponding to allocation a , the sum of a 's associated labels along any path between s and t is equal to $\pi_{N,a}(t) - \pi_{N,a}(s)$. \square

Figure 4 represents GR_N for $S_{3,2}$. On each edge, the label corresponding to $a_{1,3}$ —which denotes the change in $\pi_{N,a_{1,3}}$ due to moving along the edge—is equal to $q_{N,a_{1,3},i} \cdot \delta$ for some $i \in N$ and is exactly the negative of the label corresponding to a_2 . The the cost of the longest path (e.g., between $(c_{1,0}, c_{2,\gamma}, c_{3,0})$ and $(c_{1,\gamma}, c_{2,0}, c_{3,\gamma})$) is $3\gamma\delta$.

Now let us move to the proof of our general result.

LEMMA 4.4. *Any solution to $P(V_{N,G}^{(b)}, G)$ with $\delta > 0$ corresponds to a solution to $F(V_{N,G}^{(b)}, G)$.*

Proof. Let a solution to $P(V_{N,G}^{(b)}, G)$ for which $\delta > 0$ be given. Thus we have $\pi_{N,a}(\hat{\mathcal{V}})$ for all $\hat{\mathcal{V}} \in \mathbb{V}_{N,G}^{(b)}$. To give




Figure 4: Graph GR_N for our three-bidder, two-good example. Each node (a, b, c) denotes $(c_{1,a}, c_{2,b}, c_{3,c})$. The label corresponding to $a_{1,3}$ on directed edges from level k to $k+1$ is δ and on directed edges from level $k+1$ to k is $-\delta$, $0 \leq k < 3\gamma - 1$.

a solution to $F(V_{N,G}^{(b)}, G)$, we have to map the $\pi_{N,a}(\hat{\mathcal{V}})$'s to the allocation probabilities, winning probabilities and payments in $F(V_{N,G}^{(b)}, G)$. For all $\hat{v} \in V_{N,G}^{(b)}$, let

$$(4.2) \quad \pi_{N,a}(\hat{v}) = \pi_{N,a}(\hat{\mathcal{V}})$$

where $\hat{\mathcal{V}}_i = c_{i,s_i}$ for some $s_i \in \{0, \dots, \gamma\}$ such that $c_{i,s_i} \leq \hat{v}_i < c_{i,s_i+1}$. Also, for all $\hat{v} \in V_{N,G}^{(b)}$ and all $i \in N$ let

$$(4.3) \quad w_{N,i}(\hat{v}) = \sum_{a \in \mathbb{A}_G, a_i \geq b_i} \pi_{N,a}(\hat{v}), \text{ and}$$

$$(4.4) \quad p_{N,i}(\hat{v}) = \sum_{\substack{1 \leq s'_i \leq s_i \\ c_{i,s'_i} \leq \hat{v}_i < c_{i,s_i+1}}} c_{i,s'_i} [w_{N,i}(c_{i,s'_i}, \hat{\mathcal{V}}_{-i}) - w_{N,i}(c_{i,s'_i-1}, \hat{\mathcal{V}}_{-i})].$$

We show that the above $\pi_{N,a}$'s, $w_{N,i}$'s, $p_{N,i}$'s and δ indeed constitute a solution to $F(V_{N,G}^{(b)}, G)$.

Note that (4.3) is in fact $(F.w)$. Also note that $\delta > 0$ induces $(F.\delta)$. It is easy to see that five of the constraints in F are directly induced by (4.2), (4.3) and a set of constraints in P . Precisely, (i) $(F.a_1)$ is induced by (4.2), $(P.\pi_1)$ and $(P.\pi_2)$; (ii) $(F.a_2)$ is induced by (4.2) $(P.\pi_1)$ and $(P.\pi_3)$; (iii) $(F.max)$ is induced by (4.2) and $(P.\pi_1)$ —this simply follows the fact that any $a \notin \mathbb{M}_N$ is certainly not in \mathcal{M}_N ; (iv) $(F.w)$ is induced by (4.3); and (v) $(F.step)$ is induced by (4.2) and (4.3).

$(F.mon)$ is induced by $(P.\pi_1)$, $(P.\pi_3)$, $(P.\pi_4)$, $(P.q_1)$ – $(P.q_4)$, (4.2) and (4.3). To see this, note that we can write $w_{N,i}(\hat{v}) = \sum_{a \in \mathbb{A}_G, a_i \geq b_i} \pi_{N,a}(c_{1,s_1}, \dots, c_{n,s_n}) =$

$\sum_{a \in \mathcal{M}, a_i = b_i} (\pi_{N,a}(\hat{0}) + \sum_{\ell \in N} (q_{N,a,\ell} \cdot \delta \cdot s_\ell))$. The first equality follows from (4.2) and (4.3) and the second equality follows from $(P.\pi_1)$ and $(P.\pi_4)$. Now, if (1) $a_i = b_i, \forall a \in \mathcal{M}_N$, then, $w_{N,i}(\hat{v}) = \sum_{a \in \mathcal{M}, a_i = b_i} \pi_{N,a}(\hat{0}) = 1$. The first equality holds by $(P.q_1)$ and the second equality holds by $(P.\pi_3)$. Otherwise, let $\hat{v}' = (\hat{v}'_i, \hat{v}_{-i})$. Then $w_{N,i}(\hat{v}') - w_{N,i}(\hat{v}) = \sum_{a \in \mathcal{M}_N, a_i = b_i} (q_{N,a,i} \cdot \delta \cdot (s'_i - s_i)) = \delta \cdot (s'_i - s_i)$. The first equality holds by $(P.\pi_4)$ and the second equality holds by $(P.q_4)$. Now, if $\hat{v}_i \leq \hat{v}'_i$ then $s_i \leq s'_i$ and thus (2) $w_{N,i}(\hat{v}) \leq w_{N,i}(\hat{v}')$. Thus, by (1) and (2), we have $(F.mon)$.

$(F.cs)$ is induced by the same set of constraints as $(F.mon)$; that is, by $(P.\pi_1)$, $(P.\pi_3)$, $(P.\pi_4)$, $(P.q_1)$ – $(P.q_4)$, (4.2) and (4.3). Following the same argument as above, if $a_i = b_i, \forall a \in \mathcal{M}_N$, then, $w_{N,i}(\hat{v}) = \sum_{a \in \mathcal{M}, a_i = b_i} \pi_{N,a}(\hat{0}) = 1$. Otherwise, $w_{N,i}(c_{i,s_i+1}, \hat{v}_{-i}) - w_{N,i}(c_{i,s_i}, \hat{v}_{-i}) = \delta$. Thus we get $(F.cs)$.

$(F.sp)$ is induced by $(F.step)$ and (4.4). This is because by $(F.step)$, the integral part of $(F.sp)$ is over a discrete domain and thus we can write $(F.sp)$ as (4.4).

$(F.rm)$ is induced by (4.4) and the rest of the constraints in F . As stated above, if bidder i belongs to all $a \in \mathcal{M}_N$, then $\forall \hat{v}, w_{N,i}(\hat{v}) = 1$ and thus $p_{N,i}(\hat{v}) = 0$. Otherwise, $p_{N,i}(\hat{v}) = \sum_{1 \leq s'_i \leq s_i | c_{i,s'_i} \leq \hat{v}_i < c_{i,s_i+1}} c_{i,s'_i} \cdot \delta$. By Lemma 4.1, it is clear that dropping bidder $j \neq i$ either does not change the payment of bidder i or sets it to zero (if dropping j entails a case in which i belongs to all the allocation in the support of the mechanism). Thus $(F.rm)$ follows immediately. \square

Constraints in $P(V_{N,G}^{(b)}, G)$ are all linear or quadratic, and so our problem of identifying mechanism M_γ can be reduced to solving a set of quadratically constrained linear programs where the objective function in each is to maximize δ , and then checking each for $\delta > 0$. However, we can do even better. The next result demonstrates that this QCLP is always feasible; later, we will show how to analytically construct a solution with $\delta > 0$.

LEMMA 4.5. *Let $P(V_{N,G}^{(b)}, G)$ be given. For any given $\pi_{N,a}(\hat{0}) > 0, \forall a \in \mathcal{M}_N$, such that $\sum_{a \in \mathcal{M}_N} \pi_{N,a}(\hat{0}) = 1$, and any given $q_{N,a,i}, \forall a \in \mathcal{M}_N$, such that $(P.q_1)$ – $(P.q_5)$ are satisfied, there exists a solution to the set of inequalities in $(P.\pi_1)$ – $(P.\pi_3)$ and $(P.\pi_4)$ that sets $\delta > 0$.*

Proof. Let $\pi_{N,a}(\hat{\mathcal{V}}) = 0, \forall a \in \mathbb{A}_G - \mathcal{M}_N$. Thus we have $(P.\pi_1)$. We can write $(P.\pi_2)$ and $(P.\pi_4)$ as

$$(4.5) \quad 0 \leq \pi_{N,a}(\hat{0}) + \sum_i (q_{N,a,i} \cdot \delta \cdot s_i) \leq 1.$$

Let t_a denote $\sum_{i \in N} (q_{N,a,i} \cdot s_i)$. If $t_a = 0$ then let $\pi_{N,a} = \pi_{N,a}(\hat{0})$; by the assumption of the lemma, (4.5) holds. Otherwise, we can rewrite each equation (4.5) in which $t_a < 0$ as $\delta \leq \frac{-\pi_{N,a}(\hat{0})}{t_a}$, and $t_a > 0$ as $\delta \leq \frac{1 - \pi_{N,a}(\hat{0})}{t_a}$. Now, we have many constraints of the form $\delta \leq k$ for different nonnegative k 's. Denote the smallest k by k^* and set $\delta = k^*$. Thus (4.5) holds and furthermore, since all $\pi_{N,a}(\hat{0})$'s and all

k 's are greater than zero, $k^* = \delta$ is greater than zero. Note that if $t_a = 0$ in all constraints (4.5), any $\delta > 0$ would work.

It remains to show that $(P.\pi_3)$ holds. We can write $\sum_{a \in \mathcal{M}_G} \pi_a(\hat{\mathcal{Q}}) = \sum_{a \in \mathcal{M}_N} (\pi_{N,a}(\hat{0}) + \sum_i (q_{a,i} \cdot \delta \cdot s_i)) = \sum_{a \in \mathcal{M}_N} \pi_{N,a}(\hat{0}) + \delta \sum_{a \in \mathcal{M}_N} \sum_i (q_{N,a,i} \cdot s_i) = 1 + \delta \sum_i \sum_{a \in \mathcal{M}_N} (q_{N,a,i} \cdot s_i) = 1 + \delta \sum_i s_i \cdot (\sum_{a \in \mathcal{M}_N, a_i = \emptyset} q_{N,a,i} + \sum_{a \in \mathcal{M}_N, a_i = b_i} q_{N,a,i}) = 1 + 0 = 1$. The first equality holds by $(P.\pi_4)$, the third equality holds by the lemma's assumption, and the fifth equality holds by $(P.q_1)$, $(P.q_4)$ and $(P.q_5)$. \square

THEOREM 4.1. *For any given $\gamma \geq 0$, there exists a γ -step randomized mechanism that is strategyproof and revenue monotonic and satisfies participation, maximality and $(\gamma$ -step, $\delta)$ consumer sovereignty, for some $\delta > 0$.*

Proof. It is easy to verify that regardless of $V_{\mathbb{N},G}^{(b)}$ and G we can always generate $\pi_{N,a}(\hat{0})$'s such that $\sum_{a \in \mathcal{M}_N} \pi_{N,a}(\hat{0}) = 1$. For example, we can set $\pi_{N,a}(\hat{0}) = \frac{1}{|\mathcal{M}_N|}$. In fact, there are infinitely many such assignments of $\pi_{N,a}(\hat{0})$'s. Similarly, there exist infinitely many random assignments of the $q_{N,a,i}$'s that satisfy $(P.q_1)$ – $(P.q_5)$. Now note that except for δ , $P(V_{\mathbb{N},G}^{(b)}, G)$'s have no other variable in common. Thus, if we set δ to be the minimum of δ 's in the solutions to $P(V_{\mathbb{N},G}^{(b)}, G)$'s, the rest of the proof directly follows from Lemma 4.4 and Lemma 4.5. \square

In the example, it is clear that $3\gamma\delta \leq 1$ and thus $\delta \leq \frac{1}{3\gamma}$. Since we have complete freedom in choosing $\pi_{N,a}(\hat{0})$'s, we can set them appropriately—that is to set $\pi_{N,a_1,3}(\hat{0}) = 1/3$ and $\pi_{N,a_2}(\hat{0}) = 2/3$ —so that $\delta = \frac{1}{3\gamma}$ is feasible. As we said earlier, we described the solution of $S_{3,2}$ for $N = \mathbb{N}$ and $G = \mathbb{G}$. To fully define the mechanism we simply have to find δ for all choices for N and G and keep the smallest δ —which indeed is $\delta = \frac{1}{3\gamma}$.

Our QCLP formulation characterizes a class of randomized mechanisms that satisfy our desired properties. However, a mechanism designer may also hope to optimize some additional objective function such as social welfare or revenue. In our construction above, we have full or partial freedom to set δ , \mathcal{M}_N , $\pi_{N,a}$, and $q_{N,a,i}$. Here, we briefly discuss tuning δ ; we leave further investigations of optimization for future work.

Fixing \mathcal{M}_N 's, $\pi_{N,a}$'s and $q_{N,a,i}$'s, we obtain a strategyproof mechanism, which yields the same social welfare no matter how we set δ . This is because the social welfare depends only on the allocation and bidders' true valuations. Since fixed sets of \mathcal{M}_N 's and $\pi_{N,a}$'s always produce the same distribution over the allocation space, the social welfare is always the same in expectation regardless of δ . On the other hand, δ does affect payments; indeed, the maximum δ maximizes revenue. Furthermore, the bigger δ is, the stronger is the consumer sovereignty guarantee offered to bidders. Thus maximizing δ offers (different) benefits both to the auctioneer and to bidders.

5 A Polynomial Time Algorithm

At first glance, it may seem that constructing M_γ requires time exponential in the number of potential bidders \mathbb{N} and goods G , since P has an exponential number of variables and constraints. As the example $S_{3,2}$ in Section 4 shows, some settings are “easy to solve” because of their special bidder-bundle structure that let us circumvent the exponential nature of the problem. It turns out that, even in the general case, we can always construct \mathcal{M}_N for all $V_{\mathbb{N},G}^{(b)}$, G , and $N \subseteq \mathbb{N}$ in polynomial time in $|N|$ and $|G|$.

THEOREM 5.1. *For any given $V_{\mathbb{N},G}^{(b)}$, G and N , in time polynomial in $|N|$ and $|G|$ we can find a set of maximal allocations \mathcal{M}_N , where $\alpha_{N,G,i,s} = 0$, such that each bidder $i \in N$ belongs to at least one allocation in \mathcal{M}_N and for all $a \in \mathcal{M}_N$ and all $i \in N$, $a_i = b_i$ or $a_i = \emptyset$.*

Proof. Set $\mathcal{M}_N = \emptyset$. Randomly order all the bidders in N and mark them as “unawarded”. Then run the following greedy algorithm. (1) Set $G' = G$. Start from the top of the list and award each “unawarded” bidder i her desired bundle b_i if available, and remove b_i from G' , until there are either no more goods or no more bidders. Mark all the bidders that have awarded their desired bundle as “awarded”. (2) Start from the top of the list and award each “awarded” bidder i her desired bundle b_i if available, and remove b_i from G' , until there are either no more goods or no more bidders. (3) Add the current allocation to \mathcal{M}_N . If there is any bidder marked as “unawarded” then go to step 1. Otherwise, stop.

Steps 1-3 take $O(|N||G| \log(|G|))$ time and we have to run them at most $|N|$ times, since in each run at least one bidder is marked as “awarded”. Thus, the algorithm take time $O(|N|^2|G| \log(|G|))$ to run. We add one allocation to \mathcal{M}_N at the end of each run; thus \mathcal{M}_N is of size $O(|N|)$. \square

Unfortunately, finding a *desirable* solution among the set of feasible solutions—e.g. a solution that maximizes δ —can be complicated and dependent on the architecture of the given bidder-bundle setting, our choice of \mathcal{M}_N 's and the other parameters. However, we can construct a feasible solution in polynomial time giving a (loose) lower bound on the maximum δ that satisfies our constraints.

THEOREM 5.2. *For any given $V_{\mathbb{N},G}^{(b)}$, G and N , we can construct a γ -step randomized mechanism M_γ in time polynomial in $|N|$ and $|G|$ such that M_γ is strategyproof and revenue monotonic and satisfies participation, maximality and $(\gamma$ -step, $\delta)$ consumer sovereignty where $\delta = \frac{1}{n^2\gamma}$.*

Proof. Construct \mathcal{M}_N as in the proof of Theorem 5.1. If $|\mathcal{M}_N| = 1$, then the solution is trivial. (All participating bidders win with probability equal to one.) Otherwise, if $|\mathcal{M}_N| \geq 2$, let $q_{N,a,i} = \frac{1}{|\mathcal{M}_{N,i}|}$, $\forall a \in \mathcal{M}_N, i \in N$, where $\mathcal{M}_{N,i}$ is the set of allocations $a \in \mathcal{M}_N$ that i belongs to.

Therefore, $\frac{1}{n} \leq q_{N,a,i} \leq 1$. Let $\pi_{N,a} = \frac{1}{|\mathcal{M}_N|}$. Therefore, $\frac{1}{n} \leq \pi_{N,a} \leq \frac{1}{2}$. Consider the combination of constraints $(P.\pi_2)$ and $(P.\pi_4)$ in the form we presented in the proof of Lemma 4.5. That is $\delta \leq \frac{-\pi_{N,a}(\hat{0})}{t_a}$ if $t_a < 0$, and $\delta \leq \frac{1-\pi_{N,a}(\hat{0})}{t_a}$ if $t_a > 0$, where $t_a = \sum_{i \in N} (q_{N,a,i} \cdot s_i)$. By our choice of $\pi_{N,a}$'s and $q_{N,a,i}$'s in above, $\frac{-\pi_{N,a}(\hat{0})}{t_a} \geq \frac{1}{n^{2\gamma}}$ and $\frac{1-\pi_{N,a}(\hat{0})}{t_a} \geq \frac{1}{2n\gamma}$. Thus, $\max(\delta) \geq \min\{\frac{1}{n^{2\gamma}}, \frac{1}{2n\gamma}\} = \frac{1}{n^{2\gamma}}$. \square

6 Conclusions and Future Work

In this work, we showed that our previous impossibility result about deterministic, strategyproof, revenue monotonic CA mechanisms can be circumvented by stepwise randomized mechanisms. In future work, we intend to investigate whether our mechanisms can be extended to unknown single-minded bidders and other variations in the auction setting. We also intend to explore connections to the core and to false-name bidding, and to identify stepwise randomized mechanisms that maximize objective functions of interest. Furthermore, we aim to further investigate maximizing δ ; we conjecture that δ cannot be bounded from below by a constant.

References

- [1] G. Aggrawal and J. D. Hartline. Knapsack auctions. *SODA'07*, pages 1083–1092, 2006.
- [2] N. Andelman and Y. Mansour. A sufficient condition for truthfulness with single parameter agents. *EC'06*, 2006.
- [3] A. Archer and E. Tardos. Truthful mechanisms for one-parameter agents. *FOCS'01*, pages 482–491, 2001.
- [4] L. M. Ausubel and P. Milgrom. Ascending auctions with package bidding. *Frontiers of Theoretical Economics*, 1, 2002.
- [5] L. M. Ausubel and P. Milgrom. The lovely but lonely Vickrey auction. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*, chapter 1, pages 17–40. MIT Press, Cambridge, Massachusetts, 2006.
- [6] M. Babaioff, R. Lavi, and E. Pavlov. Single-value combinatorial auctions and implementation in undominated strategies. *SODA'06*, pages 1054–1063, 2006.
- [7] Y. Bartal, R. Gonen, and N. Nisan. Incentive compatible multi unit combinatorial auctions. *TARK'03*, pages 72–87, 2003.
- [8] A. Blum and J. D. Hartline. Near-optimal online auctions. *SODA'05*, pages 1156–1163, 2005.
- [9] P. Cramton, Y. Shoham, and R. Steinberg, editors. *Combinatorial Auctions*. MIT Press, Cambridge, Massachusetts, 2006.
- [10] R. Day and P. Milgrom. Core-selecting package auctions. *International Journal of Game Theory*, July 2007.
- [11] S. Dobzinski, N. Nisan, and M. Schapira. Truthful randomized mechanisms for combinatorial auctions. *STOC '06*, pages 644–652, 2006.
- [12] U. Feige. On maximizing welfare when utility functions are subadditive. *STOC '06*, pages 41–50, 2006.
- [13] U. Feige and J. Vondrak. Approximation algorithms for allocation problems: Improving the factor of $1 - 1/\epsilon$. *FOCS '06*, pages 667–676, 2006.
- [14] A. Goldberg and J. D. Hartline. Envy-free auctions for digital goods. *EC'03*, pages 29–35, 2003.
- [15] J. Green and J. Laffont. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica*, 45(2):427–438, 1977.
- [16] M. Haris and A. Riley. Allocation mechanisms and the design of auctions. *Econometrica*, 49(6):1477–1499, 1981.
- [17] R. Lavi, A. Mu'alem, and N. Nisan. Towards a characterization of truthful combinatorial auctions. *FOCS'03*, pages 574–583, 2003.
- [18] R. Lavi and C. Swamy. Truthful mechanism design for multi-dimensional scheduling via cycle monotonicity. *EC '07*, pages 252–261, 2007.
- [19] D. Lehmann, L. I. O'Callaghan, and Y. Shoham. Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM*, 49(5):1–26, September 2002.
- [20] A. Likhodedov and T. Sandholm. Approximating revenue-maximizing combinatorial auctions. *AAAI'05*, pages 267–274, 2005.
- [21] P. Milgrom. Incentives in core-selecting auctions. Stanford University, 2006.
- [22] A. Mu'alem and N. Nisan. Truthful approximation mechanisms for restricted combinatorial auctions: extended abstract. *AAAI'02*, pages 379–384, 2002.
- [23] R. B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- [24] N. Nisan. Introduction to mechanism design. In N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 9. Cambridge University Press, 2007.
- [25] N. Nisan and A. Ronen. Computationally feasible VCG mechanisms. *EC'00*, pages 242–252, 2000.
- [26] B. Rastegari, A. Condon, and K. Leyton-Brown. Revenue monotonicity in combinatorial auctions. *AAAI'07*, pages 122–127, 2007.
- [27] B. Rastegari, A. Condon, and K. Leyton-Brown. Revenue monotonicity in deterministic, dominant strategy combinatorial auctions. Working paper, <http://cs.ubc.ca/~baharak/research/RevenueMonotonicity.pdf>, 2008.
- [28] J. G. Riley and W. F. Samuelson. Optimal auctions. *American Economic Review*, 71(3):381–392, 1981.
- [29] A. Ronen. On approximating optimal auctions. *EC'01*, pages 11–17, 2001.
- [30] M. H. Rothkopf. Thirteen reasons why the Vickrey-Clarke-Groves process is not practical. *Operations Research*, 55(2):191–197, March-April 2007.
- [31] M. Yokoo. Pseudonymous bidding in combinatorial auctions. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*, chapter 7, pages 161–187. MIT Press, Cambridge, Massachusetts, 2006.
- [32] M. Yokoo, Y. Sakurai, and S. Matsubara. The effect of false-name bids in combinatorial auctions: New fraud in internet auctions. *Games and Economic Behavior*, 46(1):174–188, January 2004.