Pareto Optimal Budgeted Combinatorial Auctions☆

Phuong Le[1]
Analysis Group, Inc., Los Angeles

Abstract

This paper studies the possibility of implementing Pareto optimal outcomes in the combinatorial auction setting where bidders may have budget constraints. I show that when the setting involves a single good, or multiple goods but with single-minded bidders, there is a unique mechanism, called truncation VCG, that is individually rational, incentive compatible and Pareto optimal. Truncation VCG works by first truncating valuations at budgets, and then implementing standard VCG on the truncated valuations. I also provide maximal domain results, characterizing when it is possible to implement Pareto optimal outcomes and, if so, providing an implementing mechanism. Whenever there is at least one multi-minded constrained bidder and another multi-minded bidder, implementation is impossible. For any other domain, however, implementation is possible.

Keywords: Combinatorial auctions, budget constraints, Pareto optimality, single-minded
JEL Classification Code: D44, D47

1. Introduction

The progress of information technology brings ever increasing demand for telecommunications, by both end users and corporations. To meet this demand, telecommunications companies (telecoms) need to acquire more licenses for radio frequency spectrum. These licenses have been typically auctioned by the government to the telecoms. The Federal Communications Commission (FCC) spectrum auction in 1994 (Milgrom, 2000) is a prominent example. What makes spectrum auctions special is the combinatorial nature of the licenses: the value of a license depends on how it is combined with other licenses. For example, to some telecom, a license for a spectrum in California is worth \$1 million and a license for a spectrum in Nevada is also worth \$1 million, yet the combination of both licenses is worth \$5 million because the telecom can share infrastructure in the two neighboring states and reap economies of scale.

☆I thank an anonymous referee and participants in the Economic Theory seminar at Stanford University for their inputs on a previous version of this paper. I am also grateful for the comments and suggestions from two anonymous referees and the editor on the current version.

[1]Email: lp3ides@gmail.com. Website: lechiphuong.com.

Spectrum auctions are being implemented in many countries, yet many features of existing formats are not fully understood and many issues remain unresolved. One such issue is that of bidders' budget constraints: Bidders in the auction are constrained by their budgets. Continuing from the example above, the firm may be able to pay only $1 million even if it gets both licenses. The presence of budget constraints for the bidders has been detected in many auctions (Bulow et al., 2009) and is therefore a real concern.

Auctions with valuations and budgets as private information which satisfy Bayesian incentive compatibility and optimality have been studied, in the single buyer setting by Che and Gale (2000) who show that the optimal mechanism involves a menu of contracts, and in the multiple bidders setting by (Pai and Vohra, 2014) who show that the optimal auction requires "pooling" both at the top and in the middle despite the maintained assumption of a monotone hazard rate. When budgets are common knowledge, Laffont and Robert (1996) characterize the optimal auction as an all-pay auction with the appropriate reserve price. If the mechanism must be robust to the beliefs of bidders, one cannot attain optimality, constrained efficiency or ex-post efficiency, but can hope for a weaker notion of efficiency: *Pareto optimality*. Dobzinski et al. (2012) study a multi-good setting and show that when bidders are budget-constrained and budgets are private information, there is no incentive compatible Pareto optimal auction. They propose the adaptive clinching auction, a modification of the clinching auction in Ausubel (2004), and show that it satisfies Pareto optimality, individual rationality and incentive compatibility when budgets are known. They also show that the adaptive clinching auction is in fact the only such mechanism. Considering a divisible good, Hafalir et al. (2012) propose a generalization of the Vickrey auction called *Vickrey with Budgets* and show that it yields good revenue and Pareto optimality properties. Borgs et al. (2005) prove that, in the case of two buyers and two units, there is no truthful auction that allocate goods to distinct bidders. The authors also design an asymptotically revenue-maximizing truthful mechanism which may allocate only some of the items. In the general valuation environment, there is no mechanism that is Pareto optimal and incentive compatible, sometimes even with publicly known budgets (Goel et al., 2012; Dobzinski et al., 2012; Fiat et al., 2011; Lavi and May, 2012).

The current paper relaxes Pareto optimality and incentive compatibility by requiring only that they hold *almost everywhere* (in the standard Lesbegue-measure-theoretic sense), and characterizes a simple intuitive mechanism that achieves all three (potential concerns about such a relaxation are addressed in section 6, after the results have been presented). I show that, in the single-good setting, a mechanism called *truncation VCG* is individually rational, generically Pareto optimal and generically incentive compatible. Truncation VCG first truncates each bidder's valuations at his budget and then applies the usual VCG mechanism to the resulting truncated valuations, ignoring the existence of budgets. Intuitively, truncated valuations correctly capture the *willingness and ability to pay*, which provides enough information to not only attain Pareto optimality but also compute affordable payments that align incentives. I also show that any individually rational, incentive compatible and Pareto optimal mechanisms in the single-good setting must coincide with truncation VCG almost everywhere. This almost everywhere uniqueness result parallels the uniqueness of VCG in the unconstrained setting. While the uniqueness of VCG in the unconstrained setting stems from the (generically) uniqueness of the efficient allocation and the associated prices based on the taxation principle (what I call threshold prices), the uniqueness of truncation VCG is a priori not as obvious. This is because, unlike the unconstrained case, there may be multiple Pareto optimal allocations even at generic profiles. Nevertheless, incentive considerations and the threshold pricing principle require that the mechanism hold each bidder accountable for the (truncated) externality that he imposes on others. This observation allows for a complete characterization of mechanisms in this setting.

Though there are similar uniqueness results for the single-good in the literature (for example, Dobzinski et al. (2012) for two players), analysis in the single-good setting is useful in illustrating important concepts and proof techniques which are useful for more general settings. Considering multiple goods, I start with the domain where bidders are single-minded, i.e., each bidder values only a specific bundle. The results are analogous to the single-good setting: truncation VCG is essentially the unique mechanism that is individually rational, generically Pareto optimal and generically incentive compatible. The intuition is also similar to the single-good setting: as long as valuations are one-dimensional, truncation does not lose too much information. Truncation VCG can therefore be suitable for applications where the single-minded assumption is appropriate, such as auctioning of pollution rights, communication links in a tree or auto parts to buyers desiring a specific model (see Lehmann et al., 2002, and the references therein). Other potential applications are spectrum auctions where the auctioneer has sufficient information about a bidder's existing technology and wireless infrastructures to be confident that the bidder is interested in only one specific spectrum bundle.

Moving beyond single-minded valuations, I provide maximal domain results describing the domains for which Pareto optimal outcomes can be implemented, and domains for which they cannot be. When all constrained bidders are single-minded (other bidders are unconstrained and may be multi-minded, i.e., interested in multiple bundles), truncation VCG remains the unique mechanism satisfying generic individual rationality, generic incentive compatibility and Pareto optimality. As expected, the difficulty in implementation arises when constrained bidders are multi-minded. However, implementation is still possible when only one bidder is multi-minded and constrained, and other bidders are single-minded. Implementation becomes impossible when there are at least one multi-minded constrained bidder, and another multi-minded bidder (constrained or not). Collectively, these results allow one to, given any domain, determine whether implementation of Pareto optimal outcomes in dominant strategy is possible and, if so, provide an implementing mechanism.

The paper adds to the existing literature in several ways. First, instead of looking at multi-unit auctions (Dobzinski et al., 2012) or settings common for online advertising applications such as AdWords auctions (Fiat et al., 2011) or a divisible good (Bhattacharya et al., 2010), I analyze a setting which models spectrum auctions with distinct indivisible goods. Second, in my setting both valuations and budgets are private information, so the positive result in the single-minded domain stands in stark contrast against the negative results in the literature.[2] Single-minded valuations presume complementary preferences, so the results on the single-minded domain do not apply to substitutable valuations[3], but the results for the single-good setting, as well as impossibility results for the multi-minded domain, do not rely on complementarities and so hold for substitutable valuations as well.

Technically, the current paper exploits an existing characterization result in the auction literature that allows one to view any incentive compatible mechanism as an allocation rule and a payment rule which satisfy two conditions: (1) the threshold pricing condition, requiring that a bidder's payment is his threshold price, the minimum bid he must make to win his allocated bundle, and (2) the optimality condition, requiring that a bidder's allocation must be optimal for him, given his threshold prices. Pareto optimality requires that the allocation rule at unconstrained profiles must maximize total surplus, which in turn imply certain threshold prices. By construction, such threshold prices facing a bidder are independent of the bidder's type. In particular, they hold when a bidder is constrained as well. The optimality condition then determines what the allocation rule must be when bidders are constrained. In certain domains, such as the single-minded domain, it is possible to satisfy both threshold pricing and optimality, but in other domains it is not.

The paper is organized as follows. Section 2 lays out the formal environment and describes the truncation VCG mechanism. Section 3 describes the threshold pricing and optimality conditions and how they relate to incentive compatibility. Section 4 provides the results for the single-good setting, showing that truncation VCG not only has the desirable properties in this domain but also is the unique mechanism having such properties. Section 5 extends the results to the single-minded domain, shows the maximal domain results. Section 6 discusses the relaxation from full to generic incentive compatibility, and section 7 concludes.

2. Preliminaries

2.1. Setting

A seller S wants to allocate a set G of indivisible goods to a set I of bidders. Let $X(G)$ be the set of feasible allocations x, where $x = (x_1, x_2, ..., x_I)$ specifies that bidder i gets bundle x_i and must satisfy $x_i \cap x_j = \emptyset$ for all $i \neq j$. A bidder i's valuations over the bundles are summarized by a function $u_i : 2^G \to \mathbb{R}^+$. The valuation of the empty bundle is zero.

[2]In certain settings with public budgets and private valuations, e.g., in Dobzinski et al. (2012), positive results have been established. However, when both valuations and budgets are private information, negative results are obtained.

[3]Substitutable valuations, in the unconstrained setting, give rise to well-behaved demands and have important implications for existence of Walrasian equilibria as well as design of dynamic ascending price auctions (Gul and Stacchetti, 2000).

I assume each bidder only cares about his own bundle and write $u_i(x)$ to mean $u_i(x_i)$. I also assume *free-disposal*, so $u_i(y_i) \geq u_i(x_i)$ if $y_i \supseteq x_i$.

A bidder i also has a budget $b_i \in \mathbb{R}^+$ that limits how much he can afford to pay. A bidder's valuation function and budget (u_i, b_i) are private information, unknown to other bidders and the seller. A type (u_i, b_i) is *unconstrained* if $b_i \geq \max_{x_i} u_i(x_i)$. An unconstrained bidder is a bidder whose type space contains only unconstrained types. A constrained bidder is a bidder whose type space contain both constrained and unconstrained types. A profile $(u, b) = ((u_i, b_i))_{i \in I}$ describes the characteristics of all bidders. A report (u_{-i}, b_{-i}) describes the types of bidders other than i. Let $\mathbb{U}$ denote the set of all profiles, and $\mathbb{U}_i$ denote the set of i's types. It will be convenient to have $\mathbb{U}_i^u$ denote the set of i's unconstrained types.

Let $P = \{p = (p_i)_{i \in I} : p_i \in \mathbb{R}^+ \text{ for all } i\}$ be the set of payments. An outcome is a pair $(x, p) \in X(G) \times P$ that specifies that bidder i gets bundle x_i and pays p_i. Given an outcome (x, p), the payoff for bidder i is given by $v_i(x, p) = u_i(x) - p_i$ if $p_i \leq b_i$ and $-\infty$ otherwise. The seller's valuation for any bundle is assumed to be zero, and so his payoff is the total payment $v_S(x, p) = \sum_{i \in I} p_i$.

The assumption that payments are non-negative is standard in the literature, and in practice auctions generally do not pay participants. This assumption is also made in order to make the problem of finding an incentive compatible and Pareto optimal mechanism interesting. If payments can be negative, the mechanism designer can simply eliminate budget constraints, by getting the seller to make a sufficiently large transfer to each bidder so that the bidder's budget is never binding, and then implement VCG on the resulting unconstrained environment, which is known to be Pareto optimal and incentive compatible.

2.2. Mechanism

By the revelation principle, I can restrict attention to direct mechanisms. A direct mechanism elicits valuations and budgets from the bidders and then maps each profile to an outcome using a function $\phi : \mathbb{U} \to X(G) \times P$. Note that I am considering deterministic mechanisms. In a direct mechanism, each bidder's strategy space is his type space. It will be notationally convenient to split the outcome mapping ϕ into two parts: the allocation rule $\phi^a : \mathbb{U} \to X(G)$ and the payment rule $\phi^p : \mathbb{U} \to P$. Bidder i's allocation and payment at profile (u, b) shall be referred to as $\phi_i^a(u, b)$ and $\phi_i^p(u, b)$ respectively.

A mechanism is individually rational if each bidder's payoff is non-negative. This property ensures that bidders will weakly gain from participating in the mechanism. Mechanisms that are not individually rational may deter bidders from entry.

Definition 1. *A mechanism $\phi(\cdot)$ is* individually rational (IR) *if $v_i(\phi(u, b)) \geq 0$ for all i for all (u, b).*

Note that individual rationality and non-negative payments imply that losing bidders pay exactly zero.

My notion of Pareto optimality is motivated by the standard and analogous notion of Pareto optimality (which coincides with and is usually called "efficiency") in the unconstrained setting. The fact that, in the unconstrained setting, Pareto optimality is equivalent to total surplus maximization, regardless of payments, implies that the seller's welfare is considered in Pareto dominance (along with the bidders' welfare). Otherwise, one can always make the bidders weakly better off by lowering payments, implying that any Pareto optimal outcome must have zero payments. This fact also implies that the potentially Pareto-dominating comparison outcome can involve negative payments.[4] If only non-negative payments are allowed in constructing potentially Pareto-dominating comparison outcomes, then an outcome that does not maximize total surplus could be Pareto optimal. For example, consider the single-good setting with two unconstrained bidders, bidder 1 and 2, with valuations of 3 and 5 respectively. Allocating the good to bidder 1 at price of zero does not maximize total surplus, but would be considered Pareto optimal if any comparison outcome must have non-negative payments, because there is no other allocation that would make bidder 1 weakly better off. My definition of Pareto optimality considers the seller's welfare and allows for negative payments in the comparison outcome, thus coinciding with the standard notion of Pareto optimality when the bidders are unconstrained, and can be thought of as a generalization from the unconstrained setting to the constrained one.

To avoid confusion, the relaxation of non-negativity for payments of Pareto-dominating comparison outcomes is explicitly stated whenever applicable.

Definition 2. *A mechanism $\phi(\cdot)$ is* Pareto optimal at profile (u,b) *if if there is no outcome (y,q) (potentially involving negative payments) such that $v_i(y,q) \geq v_i(\phi(u,b))$ for all $i \in I \cup S$, with strict inequality for some $i \in I \cup S$. A mechanism is* Pareto optimal *if it is Pareto optimal at all profiles.*

Table 1: Illustration of Pareto optimality

	Valuations			Budget
	A	B	AB	
Bidder 1	0	3	9	3
Bidder 2	2	2	2	∞

With budget constraints, a Pareto optimal outcome need not involve the surplus-maximizing allocation. For example, consider the following environment consisting of two goods and two bidders with valuations and budgets as shown in table 1. I write $((x_1,p_1),(x_2,p_2))$ to denote the outcome where bidder i wins x_i and pays p_i, for $i=1,2$.

[4]The formal definition of an outcome allows only non-negative payments, but whenever negative payments are allowed in potentially Pareto-dominating comparison outcome, the relaxation of non-negativity will be explicitly stated.

Any outcome that involves bidder 1 getting the bundle AB maximizes total valuation and is Pareto optimal regardless of payment. The outcome $((B,3),(A,2))$ is also Pareto optimal even though it does not maximize total valuation because the seller is getting a payoff of 5, the maximum revenue possible subject to the individual rationality constraint. The outcome $((B,0),(A,0))$ shares the same allocation but is not Pareto optimal since the outcome $((AB,2)(\emptyset,-2))$ Pareto dominates it.

A mechanism is dominant strategy incentive compatible if it is in the interest of each bidder i to report his valuations and budget truthfully, regardless of the reports of other bidders.

Definition 3. *A mechanism $\phi(\cdot)$ is* incentive compatible at profile (u,b) *if for any bidder i, for all $(\hat{u}_i,\hat{b}_i)$, $v_i(\phi(u,b)) \geq v_i(\phi((\hat{u}_i,\hat{b}_i),(u_{-i},b_{-i})))$. A mechanism is* dominant strategy incentive compatible *if it is incentive compatible at all profiles.*

For the rest of the paper, I omit the qualifier "dominant strategy" and simply use "incentive compatible" to mean dominant strategy imcentive compatible.

2.2.1. VCG mechanism

The VCG mechanism stands out as the only mechanism that is individually rational, incentive compatible and Pareto optimal in the unconstrained environment (see Ausubel and Milgrom, 2006, and the references therein). It chooses the surplus maximizing allocation and charges externality-based payments. Formally, even though the bidders announce both valuations and budgets, the VCG mechanism uses only valuations to compute allocation and payments. It will be convenient to have the following notation describing the maximum total valuation and the associated maximizer(s). Given a profile (u,b), let $V^u_{\hat{I}}(x) = \sum_{i\in\hat{I}} u_i(x)$ denote the total valuation among bidders in $\hat{I}$ attained from the allocation x, and $V^u_{\hat{I}}(\hat{G}) = \max_{x\in X(\hat{G})} V^u_{\hat{I}}(x)$ denote the maximum total valuation attained among $\hat{I}$ from allocating the goods in $\hat{G}$ among the bidders in $\hat{I}$, and $x^u_{\hat{I}}(\hat{G}) = \arg\max_{x\in X(\hat{G})} V^u_{\hat{I}}(x)$ denote the associated maximizer(s). Using this notation, $V^u_I(G)$ is the maximum total valuation attained from allocating the goods in G among the bidders in I, and $V^u_{I-i}(G - x_i)$ is the maximum total valuation attained from allocating the goods in $(G - x_i)$ among the bidders in $I - i$. The VCG mechanism chooses an allocation x^* in $x^u_I(G)$ and charges payments $p_i = V^u_{I-i}(G) - V^u_{I-i}(G - x^*_i)$ which can be interpreted as the externality that i imposes by taking the bundle x^*_i, i.e., how much i reduces the total surplus among bidders in $I - i$. For each unconstrained bidder i, his payoff is $v_i(x^*_i, p_i) = u_i(x^*_i) - p_i = V^u_I(G) - V^u_{I-i}(G)$, which is non-negative and so guarantees individual rationality. Noting that the term $V^u_{I-i}(G)$ is independent of i's report, it is in i's interest to maximize the term $V^u_I(G)$, which is done through truthful reporting.

2.2.2. Truncation VCG mechanism

The VCG mechanism does not quite work in the presence of budget constraints. This is because VCG only uses valuations and so does not guarantee that the payments are affordable given the budgets and might violate individual rationality. One can try to modify the VCG mechanism to accommodate budget constraints, for example, by bounding payments from above by budgets. However, generally speaking, if allocations are based on valuations only and payments must respect budget constraints, then there are incentives to misreport (by announcing a very high valuation and a very low budget, for instance). This observation suggests that in order to satisfy incentive compatibility, both allocation and payments must be determined using both valuations and budgets. One modification of VCG that does so is to truncate each bidder's valuations at his budget, and apply VCG to the resulting truncated valuations/profile. Because VCG payments never exceed valuations, this approach also guarantees that payments are bounded above by budgets. I call this mechanism truncation VCG.

More formally, given a profile (u, b), the *truncated valuation function* for any bidder i, $\bar{u}_i : 2^G \to \mathbb{R}^+$, is defined by

$$\bar{u}_i(x_i) = \min\{u_i(x_i), b_i\} \text{ for all bundles } x_i.$$

Because truncated valuations are by construction bounded above by budgets, the budgets are never binding as long as payments do not exceed truncated valuations, which is the case when VCG is applied on truncated valuations (formally, on the now-unconstrained profile $(\bar{u}, b)$). Truncation VCG can now be formally described using the same notation of VCG mechanism, letting the truncated valuation function $\bar{u}$ replace the original valuation u. Truncation VCG chooses the allocation $x^* \in x_I^{\bar{u}}(G)$ and charges payments $p_i = V_{I-i}^{\bar{u}}(G) - V_{I-i}^{\bar{u}}(G - x_i^*)$, which can be thought of as the "truncated" externality. Because payments never exceed truncated valuations and hence never exceed budgets, the payoff for each bidder i is $v_i(x_i^*, p_i) = u_i(x_i^*) - p_i \geq \bar{u}_i(x_i^*) - p_i = V_I^{\bar{u}}(G) - V_{I-i}^{\bar{u}}(G)$. If $u_i(x_i) = \bar{u}_i(x_i)$ for all x_i, i.e., i's budget never binds, then truthful reporting is in i's interest. This is intuitive, because for an unconstrained bidder, there is no difference between truncation VCG and VCG. However, if $u_i(x_i) \neq \bar{u}_i(x_i)$ for some x_i, then truthful reporting may not be optimal for i.

Table 2: Illustration of truncation VCG

	Original Valuations			Budget
	A	B	AB	
Bidder 1	0	3	9	3
Bidder 2	2	2	2	∞

	Truncated Valuations			Budget
	A	B	AB	
Bidder 1	0	$\boxed{3}$	3	3
Bidder 2	$\boxed{2}$	2	2	∞

Example 1. *Consider the profile in table 1. The truncated valuation is shown on the right panel of table 2, with the original valuations shown on the left for ease of comparison. Truncation VCG outcome involves bidder 1 winning B and paying nothing, and bidder 2 winning A and paying nothing.*

Because I focus on deterministic mechanisms, I must resolve the issue of multiple possible outcomes when $x_I^{\bar{u}}(G)$ is multi-valued. To do so, I first index all possible outcome allocations $x \in X(G)$ and when there are multiple maximizers, choose the allocation with the lowest index.[5]

2.2.3. Generic Pareto Optimality and Generic Incentive Compatibility

The standard notion of (full) incentive compatibility turns out to be too strong in the budgeted setting. To see why, consider a simple setting with one good and two constrained bidders. A salient candidate mechanism for this setting is the second price auction, where each bidder's bid is taken to be the smaller of valuation and budget, capturing the "willingness and ability to pay." Employing this mechanism, if both bidders have valuations strictly exceeding budgets, then the bidder with the higher budget wins and pays the other bidder's budget. However, if both bidders have the same budget, then no matter who gets the good, the losing bidder always has an incentive to overstate his budget, thereby winning the good at a price equal to his budget, leading to an improvement in payoff. In other words, there is a set of "knife-edge" profiles at which incentive compatibility is not possible. I introduce the notion of *generic incentive compatibility*: a mechanism is generically incentive compatible if it is incentive compatible *almost everywhere* (*a.e.*), i.e., except for a set of profiles of measure zero.[6] A more detailed discussion of generic incentive compatibility is provided in section 6, after the main results have been presented.

Definition 4. *A mechanism is* generically incentive compatible (GIC) *if it is incentive compatible at almost all profiles.*

The notion of generic Pareto optimality will become useful when discussing truncation VCG. As it will turn out, truncation VCG is not fully Pareto optimal - there are profiles at which truncation VCG is not necessarily Pareto optimal, but the set of such profiles has measure zero.

Definition 5. *A mechanism is* generically Pareto optimal (GPO) *if it is Pareto optimal at almost all profiles.*

When discussing a mechanism, I generally write ${}^{ic}\mathbb{U}$ and ${}_{po}\mathbb{U}$ to refer to the set of profiles at which the mechanism is, respectively, incentive compatible and Pareto optimal. If a mechanism is GPO and GIC, then both ${}^{ic}\mathbb{U}$ and ${}_{po}\mathbb{U}$ contain almost all profiles, and consequently the set of profiles at which the mechanism is both incentive compatible and Pareto optimal, ${}^{ic}_{po}\mathbb{U} = {}^{ic}\mathbb{U} \cap {}_{po}\mathbb{U}$, also contains almost all profiles.

[5]Most of the results pertain to profiles where $x_I^{\bar{u}}(G)$ is a singleton, so the method of indexing does not have important implications.

[6]I use the standard Lebesgue measure on the space of profiles.

3. Sufficient and necessary conditions for generic incentive compatibility

It will be convenient to first characterize mechanisms which are individually rational, and incentive compatible at certain profiles. The notion of threshold price will be useful. Consider a mechanism $\phi(\cdot)$, with allocation rule $\phi^a(\cdot)$ and payment rule $\phi^p(\cdot)$. Given reports by other bidders (u_{-i},b_{-i}) and a bundle y_i, let $W_i(y_i,(u_{-i},b_{-i}))$ be the set of unconstrained bids that result in i winning y_i:

$$W_i(y_i,(u_{-i},b_{-i})) = \{(u_i,b_i) \in \mathbb{U}_i^u : \phi_i^a((u_i,b_i),(u_{-i},b_{-i})) = y_i\}.$$

Let $\rho(y_i,(u_{-i},b_{-i}))$ be the infimum bid on y_i that i can make to win y_i:

$$\rho_i(y_i,(u_{-i},b_{-i})) = \inf\{u_i(y_i) : (u_i,b_i) \in W_i(y_i,(u_{-i},b_{-i}))\}.$$

The threshold price facing i for a bundle x_i is defined as

$$p_i(x_i,(u_{-i},b_{-i})) = \min_{y_i \supseteq x_i} \rho_i(y_i,(u_{-i},b_{-i})).$$

Loosely speaking, if i wants to win x_i or any bundle containing x_i, then he has to make a bid of at least $p_i(x_i,(u_{-i},b_{-i}))$ on some bundle containing x_i. Conversely, if i's bid on any bundle containing x_i is strictly less than $p_i(x_i,(u_{-i},b_{-i}))$, then i does not win any bundle containing x_i.

Note that if $\phi(\cdot)$ is individually rational, then the threshold price for winning the empty bundle (i.e., losing) is zero. The notion of threshold price is useful because, coupled with incentive compatibility, it gives a lower bound on the bidder's payoff. The following lemma formalizes this idea.

Lemma 1. *If an individually rational mechanism $\phi(\cdot)$ is incentive compatible at profile (u,b), then for any bidder i, $v_i(\phi_i(u,b)) \geq \max_{x_i} v_i(x_i,p_i(x_i,(u_{-i},b_{-i})))$.*

Proof. Suppose in negation that for some bidder i, $v_i(\phi_i(u,b)) < v_i(x_i,p_i(x_i,(u_{-i},b_{-i})))$ for some bundle x_i. By individual rationality, $v_i(\phi_i(u,b)) \geq 0$, so $v_i(x_i,p_i(x_i,(u_{-i},b_{-i}))) > 0$ and $v_i(x_i,p_i(x_i,(u_{-i},b_{-i}))) = u_i(x_i) - p_i(x_i,(u_{-i},b_{-i}))$. Let $\epsilon = u_i(x_i) - p_i(x_i,(u_{-i},b_{-i})) - v_i(\phi_i(u,b)) > 0$. By definition of threshold prices, there is some unconstrained report $(\hat{u}_i,\hat{b}_i)$ with $\hat{u}_i(y_i) < p_i(x_i,(u_{-i},b_{-i})) + \epsilon$ such that i wins some bundle $y_i \supseteq x_i$ at profile $((\hat{u}_i,\hat{b}_i),(u_{-i},b_{-i}))$, yielding valuation of $u_i(y_i) \geq u_i(x_i)$. At this report, by individual rationality i pays at most $\hat{u}_i(y_i)$. So by making report $(\hat{u}_i,\hat{b}_i)$, a bidder i with true type (u_i,b_i) would have a payoff of at least $u_i(x_i) - \hat{u}_i(y_i) > u_i(x_i) - (p_i(x_i,(u_{-i},b_{-i})) + \epsilon) = v_i(\phi_i(u,b))$. So $(\hat{u}_i,\hat{b}_i)$ is a profitable misreport for bidder i, and $\phi(\cdot)$ is not incentive compatible at profile (u,b), a contradiction. $\square$

The following lemma states the sufficient conditions for a mechanism to be individually rational and incentive compatible on a set of profiles: winning bidders are charged threshold prices, and each bidder's allocation is optimal for him given the threshold prices. This is essentially the taxation principle (Wilson, 1993). Versions of this result have also been shown by Milgrom and Segal (2014); Lehmann et al. (2002); Yokoo (2003) in the unconstrained setting, and little modification is needed in the constrained setting.

Lemma 2 (Milgrom and Segal, 2014; Lehmann et al., 2002; Yokoo, 2003; Wilson, 1993)**.** *An individually rational mechanism $\phi(\cdot)$ is incentive compatible at all profiles in the set U if the following conditions hold.*

- *(Threshold pricing:) At any profile (u,b), for any bidder i, $\phi_i^p(u,b) = p_i(\phi_i^a(u,b),(u_{-i},b_{-i}))$.*
- *(Optimality:) At any profile (u,b) in U, for any bidder i, $\phi_i^a(u,b) \in \arg\max_{x_i} v_i(x_i, p_i(x_i,(u_{-i},b_{-i})))$.*

Proof. Omitted. □

Given an individually rational mechanism $\phi(\cdot)$, the report (u_{-i},b_{-i}) is called IC-typical if the set

$$^{ic}\mathbb{U}_i = \{(u_i,b_i) : \phi(\cdot) \text{ is incentive compatible at } ((u_i,b_i),(u_{-i},b_{-i}))\}$$

contains almost all i's types. The next lemma describes the outcome for bidder i at any IC-typical reports. This result is useful in characterizing generically incentive compatible mechanisms, because for such mechanisms, for any bidder i, almost all reports are IC-typical.[7] One can think of this result as the necessary conditions for incentive compatibility. Again, versions of this result have been shown by Milgrom and Segal (2014); Lehmann et al. (2002); Yokoo (2003) in the unconstrained setting. The current version is tailored to suit the weaker notion of generic incentive compatibility. The proof is essentially adopted from Yokoo (2003), with minor modifications to handle the measure-theoretic language.

Lemma 3 (Milgrom and Segal, 2014; Lehmann et al., 2002; Yokoo, 2003)**.** *Let $\phi(\cdot)$ be an individually rational mechanism. If a report (u_{-i},b_{-i}) is IC-typical then at any profile $(u,b) = ((u_i,b_i),(u_{-i},b_{-i}))$ at which $\phi(\cdot)$ is incentive compatible, the following conditions must hold.*

- *(Threshold pricing:) $\phi_i^p(u,b) = p_i(\phi_i^a(u,b),(u_{-i},b_{-i}))$.*
- *(Optimality:) $\phi_i^a(u,b) \in \arg\max_{x_i} v_i(x_i, p_i(x_i,(u_{-i},b_{-i})))$.*

[7] Suppose otherwise, then there must be a set of strictly positive measure of reports $\hat{U}_{-i}$ such that at any report (u_{-i},b_{-i}) in $\hat{U}_{-i}$, there is a set of strictly positive measure of i's types $\hat{U}_i$ such that for any (u_i,b_i) in $\hat{U}_i$, $\phi(\cdot)$ is not incentive compatible at $(u,b) = ((u_i,b_i),(u_{-i},b_{-i}))$. But this means there is a set of profiles with strictly positive measure at which $\phi(\cdot)$ is not incentive compatible, contradicting generic incentive compatibility.

Note that unlike the sufficient conditions version, the above conditions need to hold only at profiles where $\phi(\cdot)$ is incentive compatible. The lemma is stated without any restriction on the bidders' types, and holds for all domains.

Proof. Suppose that the report (u_{-i},b_{-i}) is IC-typical, associated with the set ${}^{ic}\mathbb{U}_i$. Consider any profile (u,b) at which $\phi(\cdot)$ is incentive compatible. To show threshold pricing, suppose that bidder i is allocated bundle x_i (which can be the empty bundle) at this profile. By lemma 1, i's payoff is bounded below by $v_i(x_i,p_i(x_i,(u_{-i},b_{-i})))$, so $\phi_i^p(u,b) \leq p_i(x_i,(u_{-i},b_{-i}))$. Suppose $\phi_i^p(u,b) < p_i(x_i,(u_{-i},b_{-i}))$. Because ${}^{ic}\mathbb{U}_i$ contains almost all i's types, there is some unconstrained type $(\hat{u}_i,\hat{b}_i) \in {}^{ic}\mathbb{U}_i$ such that for some small $\epsilon > 0$,

$$\begin{cases} \hat{u}_i(y_i) \in (c_i - \epsilon, c_i + \epsilon) & \text{for all } y_i \supseteq x_i; \\ \hat{u}_i(y_i) < c_i - \epsilon - \phi_i^p(u,b) & \text{for all } y_i \not\supseteq x_i. \end{cases}$$

for some c_i satisfying $\phi_i^p(u,b) < c_i - \epsilon < c_i + \epsilon < p_i(x_i,(u_{-i},b_{-i}))$. The choice of $(\hat{u}_i,\hat{b}_i)$ ensures that his outcome yields a payoff less than $c_i - \epsilon - \phi_i^p(u,b)$. To see this, note that he does not win any bundle containing x_i by the definition of threshold prices. If he wins a bundle not containing x_i, then his valuation on that bundle is less than $c_i - \epsilon - \phi_i^p(u,b)$, so his payoff must also be less than $c_i - \epsilon - \phi_i^p(u,b)$ (because payments are non-negative).

Because $(\hat{u}_i,\hat{b}_i)$ is in ${}^{ic}\mathbb{U}_i$, $\phi(\cdot)$ must be incentive compatible at $((\hat{u}_i,\hat{b}_i),(u_{-i},b_{-i}))$. However, the type $(\hat{u}_i,\hat{b}_i)$ can deviate to (u_i,b_i) and, by assumption, win x_i at price $\phi_i^p(u,b)$, getting a payoff of at least $c_i - \epsilon - \phi_i^p(u,b)$, thereby improving his payoff. This contradicts incentive compatibility at $((\hat{u}_i,\hat{b}_i),(u_{-i},b_{-i}))$. Therefore, $\phi_i^p(u,b) = p_i(x_i,(u_{-i},b_{-i}))$.

To show optimality, let $v_i^* = \max_{x_i} v_i(x_i,p_i(x_i,(u_{-i},b_{-i})))$. By lemma 1, i's payoff is bounded below by v_i^*. By threshold pricing, if i wins x_i then he pays $p_i(x_i,(u_{-i},b_{-i}))$, so his payoff is bounded above by v_i^*. Therefore, i's payoff is exactly v_i^*. In order to attain this payoff, i's allocation must be in $\arg\max_{x_i} v_i(x_i,p_i(x_i,(u_{-i},b_{-i})))$. □

Lemma 3 allows one to think of any generically individually rational and incentive compatible mechanism as an allocation rule and a payment rule, the latter being the threshold prices associated with the former. The optimality condition expresses the relationship between the allocation rule and threshold prices.

4. Results for the single-good setting

In this section I restrict attention to the single-good setting, and characterize mechanisms that are IR, GPO and GIC in this setting. Even though the focus on the single-good setting takes "combinatorial" away from the auction, it has its advantages. First, the simple and canonical single-good setting is easy to understand, simplifying the notation, the results and the proofs. Second, the characterization results in this setting can be compared to the unbudgeted single-good setting to see the effect of budget constraints on the set of satisfactory mechanisms. Lastly, many of the results hold when “combinatorial” is put back, such as in the case of single-minded bidders.

In the single-good setting, there is only one good up for sale, and each bidder has a valuation for this good, and a budget. It will be convenient to simplify the notation. Bidder i's valuation is denoted by a real number u_i, his budget is still denoted by b_i, and his truncated valuation for the good is denoted by $\bar{u}_i = \min(u_i, b_i)$. Because there is only one good, the allocation is sufficiently specified through the identity of the winner. I write $x(u,b) = i$, or, equivalently, $x_i(u,b) = 1$, to mean that the allocation at profile (u,b) involves giving the good to bidder i. For the single-good setting, truncation VCG simplifies to giving the good to the bidder with the highest truncated valuation $x^* = \arg\max_{k \in I} \bar{u}_k = i^*$ (when there are multiple such bidders, break ties arbitrarily), and charging him the second highest truncated valuation $p_{i^*} = \max_{k \in I - i^*} \bar{u}_k$. Losing bidders pay zero.

As shown earlier, in the presence of budget constraints there are generally many allocations that are Pareto optimal. Furthermore, unlike the unconstrained case, whether an outcome is Pareto optimal depends on not just the allocation but also the payments. There are at least two exceptions, however. First, any outcome involving the valuation-maximizing allocation is Pareto optimal regardless of payments. Second, provided that there is only one bidder with the maximum truncated valuation, any outcome in which this maximal bidder wins the good is Pareto optimal. The following lemma formalizes the second observation. Let $\mathbb{U}_t$ be the set of profiles at which only one bidder has the highest truncated valuation, called the *maximal* bidder.

Lemma 4. *At any profile (u,b) in $\mathbb{U}_t$, any individually rational outcome (x,p) such that $x = \arg\max_{i \in I} \bar{u}_i$ is Pareto optimal.*

Proof. Let $(u,b) \in \mathbb{U}_t$ be given, and consider any outcome (x,p) in which the maximal bidder, denoted by m, wins the object at price p_m, and other bidders lose and pay zero. Suppose in negation that another outcome (y,q) (potentially involving negative payments) Pareto dominates (x,p). There are two possible cases for (y,q): (1) bidder m still wins the object, and (2) bidder m no longer wins the object. In case (1), because allocation is unchanged ($x = y = m$), Pareto dominance implies that for any bidder $i \in I$,

$$q_i \leq p_i. \tag{1}$$

But this implies that the payoff for the seller at (y,q) is weakly smaller than his payoff at (x,p):

$$\sum_{i \in I} q_i \leq \sum_{i \in I} p_i. \tag{2}$$

By Pareto dominance, the seller must not be worse off, so inequality 2 must be an equality, which means that the inequalities in 1 must all be equalities. This in turn implies that $(x,p) = (y,q)$, contradicting the assumption that (y,q) Pareto dominates (x,p).

In case (2), suppose that at allocation y bidder k wins the good instead of bidder m.

Because $(u,b) \in \mathbb{U}_t$,

$$\bar{u}_m > \bar{u}_k. \tag{3}$$

Because (x,p) is individually rational and is Pareto-dominated by (y,q), (y,q) must also be individually rational, so k can pay at most

$$q_k \leq \bar{u}_k. \tag{4}$$

For any bidder i other than m and k, who does not win at both (x,p) and (y,q), Pareto dominance and the fact that these bidders lose at (x,p) imply, respectively, that $q_i \leq p_i$ and $p_i = 0$, which combine to yield

$$q_i \leq 0. \tag{5}$$

Consider the outcome (x,p). At this outcome, the losing bidders pay exactly zero, so the payoffs for bidder m and the seller are, respectively, $u_m - p_m \geq \bar{u}_m - p_m$, and p_m, so the total payoff to the seller and m is weakly greater than $\bar{u}_m$. By Pareto dominance, the outcome (y,q) must allocate a total valuation of at least $\bar{u}_m$ to the seller and m. This total valuation must come in the form of payments from k and other remaining bidders in $(I - m - k)$. By the inequalities 4 and 5, this total payment is at most $\bar{u}_k$, which is strictly less than $\bar{u}_m$ (by inequality 3). Therefore, (y,q) cannot Pareto-dominate (x,p). □

Note that for profiles with multiple maximal bidders, there are outcomes involving giving the good to a maximal bidder, but which are not Pareto optimal. For example, consider the setting with two bidders, 1 and 2, whose types are $(u_1,b_1) = (2,2)$ and $(u_2,b_2) = (9,2)$ (same budgets, but one bidder is unconstrained and the other is constrained). The outcome in which bidder 1 wins the good at price 2 is not Pareto optimal, because it is dominated by the outcome where bidder 2 wins the good at price 2.

I can now state and prove the first main result.

Theorem 1. *In the single-good setting, truncation VCG is IR, GPO and GIC.*

Proof. Just like standard VCG in the unconstrained setting never makes a bidder pay more than his valuation, truncation VCG never makes a bidder pay more than his truncated valuation and is therefore individually rational. Recall the set $\mathbb{U}_t$ of profiles where there is only one maximal bidder ($\arg\max_{i \in I} \bar{u}_i$ is a singleton). The complement of $\mathbb{U}_t$ consists of profiles where $\arg\max_{i \in I} \bar{u}_i$ is multi-valued, and has measure zero, so $\mathbb{U}_t$ contains almost all profiles. I now argue that truncation VCG is incentive compatible and Pareto optimal at profiles in $\mathbb{U}_t$.

Incentive compatibility on the set $\mathbb{U}_t$ is shown by verifying the conditions of threshold pricing and optimality in lemma 2. Because truncation VCG gives the good to the bidder with the highest truncated valuation, it induces threshold prices facing a bidder i:

$$p_i(x_i,(u_{-i},b_{-i})) = \begin{cases} \max_{k \in I-i} \bar{u}_k & \text{if } x_i = 1; \\ 0 & \text{otherwise.} \end{cases}$$

Essentially, a bidder must bid at least the maximum truncated valuation among other bidders in order to win. This threshold price coincides with the truncation VCG payment (the winner pays the second highest truncated valuation and losing bidders pay zero), so truncation VCG satisfies threshold pricing. To show optimality, first consider the maximal bidder m. Given that his valuation is, by definition, strictly greater than his threshold price (which is the second highest truncated valuation), it is optimal for him to win the object and pay the price. For any other bidder $i \neq m$, his threshold price is $\bar{u}_m$, which is strictly greater than i's truncated valuation, so it is optimal for him not to win the object.

Pareto optimality on $\mathbb{U}_t$ follows from lemma 4, noting that truncation VCG, given truthful reporting, allocate the object to the maximal bidder and is individually rational. □

Note that truncation VCG is not fully incentive compatible, i.e., there is a set of profiles (of measure zero) with multiple maximal bidders at which truncation VCG creates incentives to misreport. For example, consider the setting with two bidders, 1 and 2, with types $(u_1,b_1) = (9,2)$ and $(u_,b_2) = (5,2)$. Truncation VCG assigns the good to either bidder, say bidder 1, at price 2, but bidder 2 has the incentive to report a higher truncated valuation (such as $(\hat{u}_2,\hat{b}_2) = (5,3)$) and win the good at price 2.

Theorem 1 shows that the single-good setting allows for the existence of an IR, GPO and GIC mechanism, namely truncation VCG. It is natural to inquire whether other such mechanisms exist for this setting. I provide an answer in the negative: truncation VCG is the only such mechanism. The main idea of the proof is to induct on the number of constrained bids. A *constrained bid* is a bid (u_i,b_i) such that $b_i < u_i$. If there is no constrained bid, then Pareto optimality coincides with surplus maximization and requires that the bidder with the highest valuation wins the good. This allocation rule then determines threshold prices for any bidder, *at any report containing no constrained bids*, to be the (truncation) VCG threshold prices. Now consider profiles with one constrained bid and the bidder making the only constrained bid, say bidder i. Because the report by other bidders contains no constrained bid, bidder i's threshold prices are already known from the previous step. The optimality condition then determines i's allocation. The allocation for the remaining bidders is determined by Pareto optimality. Now I have the allocation rule for profiles containing at most one constrained bid and, with it, the associated threshold prices for any bidder, for any report containing at most one constrained bid. I can now use the same reasoning to determine allocation at profiles with at most two constrained bids, and so on, until allocation for all profiles is determined. Effectively, Pareto optimality and incentive compatibility allows me to "extend" the allocation rule from the space of profiles with no constrained bids to the space of profiles with at most one constrained bids, and then to the space of profiles with at most two constrained bids, and so on, until completion.

Theorem 2. *In the single-good setting, any mechanism that is IR, GPO and GIC must coincide with truncation VCG a.e.*

Sketch of proof. Consider any mechanism $\phi(\cdot)$ that is IR, GPO, and GIC. Let $\mathbb{U}^k = \{(u,b) \in \mathbb{U} : |\{i : b_i < u_i\}| \leq k\}$ be the set of profiles containing at most k constrained bids. The proof first establishes, through induction, that the allocation rule $\phi^a(\cdot)$ coincides with truncation VCG's allocation rule *a.e.* through three claims. I provide here only the sketch of the steps involved in order to highlight the underlying structure of the proof and its key components. For a formal proof, please refer to the appendix.[8]

Claim 1 (Base case)**.** *At almost all profiles (u,b) in $\mathbb{U}^0$, $\phi^a(u,b) = \arg\max_{i\in I} \bar{u}_i$.*

Sketch of proof of claim 1. Note that at profiles in $\mathbb{U}^0$, no bidder is constrained. Pareto optimality and perfect transferability of utility pin down the allocation rule to be the one that maximizes total surplus, i.e., the bidder(s) with the highest valuation wins. Because bidders are not constrained, such bidders are the maximal bidder(s). If we restrict attention to profiles with only one maximal bidder, i.e., the set $(\mathbb{U}^0 \cap \mathbb{U}_t)$ which contains almost all profiles, then the allocation is completely and uniquely determined by $\arg\max_{i\in I} \bar{u}_i$.

Claim 2 (Inductive step - threshold prices)**.** *Suppose that at almost all profiles in $\mathbb{U}^k$, $\phi^a(u,b) = \arg\max_{i\in I} \bar{u}_i$. Then at almost all profiles $(u,b) \in \mathbb{U}^k$, the threshold prices for any bidder i at the report (u_{-i},b_{-i}) are given by:*

$$p_i(x_i,(u_{-i},b_{-i})) = \begin{cases} \max_{k\in I-i} \bar{u}_k & \textit{if } x_i = 1; \\ 0 & \textit{otherwise.} \end{cases}$$

Note that this threshold price formula applies to almost all reports (u_{-i},b_{-i}) containing at most k constrained bids.

Sketch of proof of claim 2. It is easy to see that if the allocation rule $\phi^a(u,b) = \arg\max_{i\in I} \bar{u}_i$ (the maximal bidder(s) always win) actually held *everywhere* in $\mathbb{U}^k$, then the threshold price formula (one must bid weakly higher than all other bidders to win) would hold everywhere in $\mathbb{U}^k$. The "almost all" qualification on the threshold price formula comes from its counterpart in the allocation rule.

Claim 3 (Inductive step - allocation)**.** *Suppose that at almost all profiles in $\mathbb{U}^k$, $\phi^a(u,b) = \arg\max_{i\in I} \bar{u}_i$. Then at almost all profiles $(u,b) \in \mathbb{U}^{k+1}$, the allocation is $\phi^a(u,b) = \arg\max_{i\in I} \bar{u}_i$.*

Sketch of proof of claim 3. Consider a profile (u,b) in $\mathbb{U}^{k+1}$, and a bidder i who is constrained at this profile. Because i is constrained, the report (u_{-i},b_{-i}) now contains at most k constrained bids, so the threshold price formula in claim 2 applies to i. This argument works for all constrained bidders, so this threshold price formula applies to all constrained bidders at (u,b).

Assume for now that there is a unique maximal bidder $m = \arg\max_{i\in I} \bar{u}_i$. By definition, $u_m \geq \bar{u}_m > \max_{k\in I-m} \bar{u}_k$. Consider two cases: (1) m is constrained and (2) m is not. If m is

[8] To avoid repetition, the appendix includes the proof of theorem 3 which, though formally written for the more general single-minded domain, applies verbatim as proof of theorem 2 as well.

a constrained bidder, then by the threshold price formula in claim 2, $\max_{k\in I-m}\bar{u}_k$ is his threshold price for winning, so by the necessary optimality condition (lemma 3), m must win. If m is an unconstrained bidder, then for any constrained bidder i, his threshold price for winning is $\max_{k\in I-i}\bar{u}_k = \bar{u}_m > \bar{u}_i$, so the optimality condition implies i must lose. So all constrained bidders lose, and the good is allocated to some unconstrained bidder (if at all). By Pareto optimality and perfect transferability of utility among unconstrained bidders, the good must be given to the bidder with the highest valuation, who is m. Therefore, $\phi^a(u,b) = m = \arg\max_{i\in I}\bar{u}_i$.

The "almost all" qualification comes from restricting attention to profiles which satisfy the conditions for (1) the threshold price formula, (2) uniqueness of maximal bidder, (3) threshold pricing and optimality of lemma 3.

Claims 1 through 3 establish that the allocation rule of $\phi(\cdot)$ coincides with truncation VCG *a.e.* By the necessary threshold pricing condition in lemma 3, $\phi(\cdot)$ charges threshold prices which, as shown in claim 2, coincides with truncation VCG payment. Therefore, $\phi(\cdot)$ coincides with truncation VCG *a.e.* □

5. Maximal domain results

5.1. *Results for the single-minded domain*

I begin the analyses on maximal domains with the single-minded domain, in which there are multiple goods and each bidder is interested in one bundle of goods only. Studying the single-minded domain is useful for two reasons. First, the single-minded domain is of independent interest and is applicable in certain situations mentioned in the introduction. Second, the single-minded domain serves as a bridge between the single-good setting and the more general combinatorial setting. Even though valuations in the single-minded domain are combinatorial, they remain one-dimensional in the sense that they are summarized by a real number and a bundle, similar to the single-good setting. As a consequence, the results from the single-good setting carry over to the single-minded domain, and can be relied upon for analysis in the more general combinatorial setting.

Definition 6. *A bidder i is* single-minded *if there is a bundle $\bar{x}_i$ such that for any bundle y_i,*

$$u_i(y_i) = \begin{cases} u_i(\bar{x}_i) & \text{if } y_i \supseteq \bar{x}_i; \\ 0 & \text{otherwise.} \end{cases}$$

The characteristics of a single-minded bidder i is summarized by $(u_i,b_i) = (\bar{x}_i,c_i,b_i)$ where $\bar{x}_i$ denotes the bundle of interest, c_i denotes $u_i(\bar{x}_i)$, and b_i denotes the budget constraint. The profile in table 3 is an example of single-minded valuations. Bidder 1's single-minded bid is summarized by $(\{A\},7,3)$. It is natural and convenient to restrict attention to mechanisms with *exact* allocation rules where a single-minded bidder's allocation is either the empty set (he is losing) or exactly his bundle of interest. The exactness restriction is without loss of generality and the results in the current paper still hold with the appropriate but not substantive modifications.[9]

Table 3: Illustration of single-mindedness

	Valuations			Budget
	A	B	AB	
Bidder 1	7	0	7	3
Bidder 2	0	6	6	∞
Bidder 3	0	0	10	5

The first result in the single-minded domain that parallels its counterpart in the single-good setting pertains to the Pareto optimality of outcomes, which involve allocations which maximize truncated valuations. Analogous to the single-good setting, let $\mathbb{U}_t$ be the set of profiles at which the truncated valuation maximizer $x_I^{\bar{u}}(G) = \arg\max_{x \in X(G)} V_I^{\bar{u}}(x)$ is unique. Note that this result does not hold in the multi-minded domain, as shown in the discussion following table 1.

Lemma 5. *At any profile (u, b) in $\mathbb{U}_t$, any individually rational outcome (x, p) such that $x = x_I^{\bar{u}}(G)$ is Pareto optimal.*

Proof. Omitted. □

Similar to the single-good setting, truncation VCG is IR, GPO and GIC in the single-minded domain. Moreover, it is essentially the unique mechanism in this setting with these properties.

Theorem 3. *In the single-minded domain, truncation VCG is IR, GPO and GIC. Moreover, any mechanism that is IR, GPO and GIC in the single-minded domain must coincide with truncation VCG a.e.*

Proof. Omitted. □

5.2. Maximal domain results

The results thus far answer two questions: (1) whether Pareto optimal outcomes can be implemented in the single-good and the single-minded domains and (2) the ways in which such implementation can be done. This section addresses a maximal domain question: In what domains can Pareto optimal outcomes be implemented in dominant strategy? Alternatively, how big must a domain get before such implementation is impossible?

[9]Intuitively, winning a bundle strictly contained by the bundle of interest gets a valuation of zero and is equivalent to losing, and winning a bundle strictly containing the bundle of interest does not increase the valuation, yet may incur higher payment. Therefore, there are effectively two relevant alternatives for a bidder: (a) lose, or (b) win the bundle of interest.

Table 4: Type space of a bidder

Types	Single-Minded	Multi-Minded
Constrained	D_1	D_2
Unconstrained	D_3	D_4

I refer to bidders who are not single-minded, i.e., who may be interested in multiple bundles, as multi-minded. It is helpful to clarify my description of various domains. A domain describes the type space for each bidder, who can belong to one of the four type spaces listed in table 4. Thus far I have considered the domain in which *every* bidder has type space D_1 (the single-good setting can be thought of as a special case of the single-minded domain) and shown that truncation VCG is essentially the unique mechanism that is IR, GPO and GIC in this domain. It is known that (1) in the domain in which every bidder is unconstrained (single- or multi-minded), VCG is the unique mechanism satisfying individual rationality, incentive compatibility and Pareto optimality and (2) in the domain in which every bidder is constrained and multi-minded (D_2), there is no mechanism satisfying individual rationality, incentive compatibility and Pareto optimality. The results in this section pertain to domains in which different bidders may have different type spaces.

In the single-minded domain, truncation of valuations does not lose too much information because a bidder is essentially one-dimensional: truncated valuations correctly capture the *willingness and ability to pay*, which in turn provides enough information to not only attain Pareto optimality but also compute affordable payments that align incentives. However, when a constrained bidder is multi-minded, i.e., interested in multiple bundles, truncation of valuations loses valuable information, namely how much he prefers one bundle over another. As shown in example 1, this loss of information prevents truncation VCG from attaining either Pareto optimality or incentive compatibility. This suggests that the difficulty in implementation lies with constrained multi-minded bidders. The following result makes this precise.

Theorem 4. *In the domain in which constrained bidders are single-minded, truncation VCG is IR, GPO and GIC. Moreover, any mechanism that is IR, GPO and GIC in this domain must coincide with truncation VCG a.e.*

The formal proof is deferred to the appendix. For some intuition, consider truncation VCG. For any unconstrained bidder, truncation VCG is effectively just VCG on his true valuations and others' (potentially) modified valuations. Therefore, it is incentive compatible for him, even if he is multi-minded. For a constrained but single-minded bidder, truncation VCG is GIC for the same reasons it is GIC in the single-minded domain. The argument showing Pareto optimality is very similar to the proof of lemma 5. The argument for why any mechanism that is IR, GPO and GIC in this domain must coincide with truncation VCG *a.e.* is practically identical to the inductive proof of the "necessary" part of theorem 3.

Theorem 4 implies implementation is impossible only when *constrained bidders are multi-minded*, i.e., there must be some bidder with type space D_2. However, this impossibility is not guaranteed when *only one constrained bidder is multi-minded*. In fact, in such a domain, implementation is still possible, via a variant of truncation VCG, as shown the next theorem. But first, the following lemma would be useful. Let $x^{\bar{u}}_{C,j}(\hat{G})$ and $x^{\bar{u}}_{C,-j}(\hat{G})$ denote the allocation for j and bidders other than j, respectively, at the unique[10] maximizer $x^{\bar{u}}_{C}(\hat{G})$.

Lemma 6. *Let $(\hat{u}_j, \hat{b}_j)$ and (u_j, b_j) be two single-minded reports with the same bundle of interest $\bar{x}_j$ such that $\bar{u}_j(\bar{x}_j) - \bar{\hat{u}}_j(\bar{x}_j) = \delta > 0$. Let C be any set of bidders and $\hat{G}$ be any set of goods. Let $\bar{\hat{u}}$ and $\bar{u}$ be the truncated valuations associated with profile $((\hat{u}_j, \hat{b}_j), (u_{-j}, b_{-j}))$ and profile $((u_j, b_j), (u_{-j}, b_{-j}))$ respectively. Then the following statements are true.*

1. *If $x^{\bar{u}}_{C,j}(\hat{G}) = \emptyset$, then $V^{\bar{u}}_{C}(\hat{G}) = V^{\bar{\hat{u}}}_{C}(\hat{G})$.*
2. *If $x^{\bar{\hat{u}}}_{C,j}(\hat{G}) = \bar{x}_j$, then $V^{\bar{u}}_{C}(\hat{G}) = V^{\bar{\hat{u}}}_{C}(\hat{G}) + \delta$.*
3. *$V^{\bar{\hat{u}}}_{C}(\hat{G}) + \delta \geq V^{\bar{u}}_{C}(\hat{G}) \geq V^{\bar{\hat{u}}}_{C}(\hat{G})$.*

Intuitively, if the truncated valuation maximizer does not involve allocating any good to j, then a lower bid from j (from $\bar{u}_j$ to $\bar{\hat{u}}_j$) does not change total truncated valuation (because the truncated valuation maximizer is unchanged). Conversely, if truncated valuation maximizer involves allocating some good to j, then a higher bid from j (from $\bar{\hat{u}}_j$ to $\bar{u}_j$) increases total valuation by the same quantity δ (again, because the truncated valuation maximizer is unchanged). Finally, an increase in j's bid cannot decrease total truncated valuation, and cannot increase it by more than the increase δ in bid. The proof of lemma 6 is straightforward, relies on the principle of optimality and is deferred to the appendix.

Theorem 5. *In the domain in which only one bidder is multi-minded (constrained or not), there exists a mechanism that is IR, GPO and GIC.*

Proof. The proof first describes a candidate mechanism in step 1, then shows that it is IR, and GIC in step 2, by verifying threshold pricing and optimality conditions. GPO is shown in the appendix.
Step 1: A candidate mechanism.

Let the only multi-minded bidder be bidder i. Consider the allocation rule $\phi^a(\cdot)$ described by the following procedure.

1. Let $\hat{p}_i(x_i(u_{-i}, b_{-i})) = V^{\bar{u}}_{I-i}(G) - V^{\bar{u}}_{I-i}(G - x_i)$.
2. Allocation for bidder i is: $x^*_i = \arg\max_{x_i} v_i(x_i, \hat{p}_i(x_i, (u_{-i}, b_{-i})))$.

[10]Uniqueness is assumed only for notational convenience and is not necessary for the results to hold.

3. Allocation for other bidders is: $x^*_{I-i} = x^{\bar{u}}_{I-i}(G - x^*_i)$.

To get some intuition behind the construction of this allocation rule, consider the multi-minded bidder i and suppose for the sake of the argument that i is actually single-minded, with bundle of interest x_i. By theorem 3, for any mechanism that is IR, GPO and GIC, at most profiles (u,b), the allocation rule must be $x^{\bar{u}}_I(G)$. The associated threshold price function for bidder i is $\hat{p}_i(x_i,(u_{-i},b_{-i})) = V^{\bar{u}}_{I-i}(G) - V^{\bar{u}}_{I-i}(G - x_i)$. In order to be incentive compatible, the optimality condition in lemma 3 requires that when i reports a multi-minded valuation, his allocation must be the optimal bundle for him given his threshold prices, i.e. his allocation must be $x^*_i = \arg\max_{x_i} u_i(x_i) - \hat{p}_i(x_i,(u_{-i},b_{-i}))$. This motivates steps 1 and 2 of the procedure. The allocation for other bidders is to maximize truncated valuation from the remaining goods.

Consider the payment rule $\phi^p(\cdot)$ that charges each bidder the threshold price associated with the allocation rule $\phi^a(\cdot)$, and the mechanism $\phi(\cdot)$ defined by $\phi(\cdot) = (\phi^a(\cdot),\phi^p(\cdot))$. Consider the set of profiles $\mathbb{U}_\phi$ for which the allocation rule $\phi^a(\cdot)$ yields unique allocations. It is easy to see that $\mathbb{U}_\phi$ contains almost all profiles. I now argue that the mechanism $\phi(\cdot)$ is individually rational, and generically incentive compatible by verifying the conditions of threshold pricing, and optimality on the set $\mathbb{U}_\phi$.

Setp 2: Individual rationality and optimality of $\phi(\cdot)$.

Threshold pricing holds by construction of $\phi(\cdot)$. Consider a profile (u,b) in $\mathbb{U}_\phi$. It is easy to see that bidder i's threshold price is simply $p_i(x_i,(u_{-i},b_{-i})) = V^{\bar{u}}_{I-i}(G) - V^{\bar{u}}_{I-i}(G - x_i) = \hat{p}_i(x_i,(u_{-i},b_{-i}))$.[11] By step 2, the outcome is both individually rational and optimal for bidder i. For any other bidder $j \neq i$, because the mechanism operates through truncated valuation, threshold pricing ensures that payments never exceed truncated valuation, so individual rationality is guaranteed. Now I show optimality for bidder j. By definition of $\mathbb{U}_\phi$, either (a) $\bar{u}_j(\bar{x}_j) > p_j(\bar{x}_j,(u_{-j},b_{-j}))$ or (b) $\bar{u}_j(\bar{x}_j) < p_j(\bar{x}_j,(u_{-j},b_{-j}))$. I show that in case (a) bidder j wins, and in case (b) bidder j loses.

Suppose that $\bar{u}_j(\bar{x}_j) > p_j(\bar{x}_j,(u_{-j},b_{-j}))$. By the definition of threshold price, there is some unconstrained bid $(\hat{u}_j,\hat{b}_j)$ such that $\bar{u}_j(\bar{x}_j) > \bar{\hat{u}}_j(\bar{x}_j) > p_j(\bar{x}_j,(u_{-j},b_{-j}))$ and j wins $\bar{x}_j$ at profile $((\hat{u}_j,\hat{b}_j),(u_{-j},b_{-j}))$. Let $\bar{\hat{u}}$ and $\bar{u}$ be the truncated valuations associated with profile $((u_{-j},b_{-j}),(\hat{u}_j,\hat{b}_j))$ and profile $((u_{-j},b_{-j}),(u_j,b_j))$ respectively.

Now consider the bundle y_i that is optimal for bidder i via steps 1 and 2 of the allocation rule, at the profile (u,b), and in particular, the truncated valuation maximizer of the remaining goods, $x^{\bar{u}}_{I-i}(G - y_i)$. I show that $x^{\bar{u}}_{I-i,j}(G - y_i) = \bar{x}_j$. Suppose otherwise

[11]If bidder i makes a single-minded bid of $\hat{p}_i(x_i,(u_{-i},b_{-i})) + \epsilon$ on x_i, then he wins x_i according to the allocation rule, so his threshold price is at most $\hat{p}_i(x_i,(u_{-i},b_{-i}))$. For any bundle y_i containing x_i, $V^{\bar{u}}_{I-i}(G - y_i) \leq V^{\bar{u}}_{I-i}(G - x_i)$ which means $\hat{p}_i(y_i,(u_{-i},b_{-i})) \geq \hat{p}_i(x_i,(u_{-i},b_{-i}))$. So if i bids less than $\hat{p}_i(x_i,(u_{-i},b_{-i}))$ on all bundles y_i containing x_i, step 2 of the allocation rule implies he does not win any such bundle y_i containing x_i. Therefore, his threshold price is $\hat{p}_i(x_i,(u_{-i},b_{-i}))$.

that $x^{\bar{u}}_{I-i,j}(G-y_i)=\emptyset$, and consider the threshold price of y_i for i at profile $(\hat{u},\hat{b})$.

$$\begin{aligned} & p_i(y_i,(\hat{u}_{-i},\hat{b}_{-i})) = V^{\bar{\hat{u}}}_{I-i}(G) - V^{\bar{\hat{u}}}_{I-i}(G-y_i) && (6)\\ \to\ & p_i(y_i,(\hat{u}_{-i},\hat{b}_{-i})) \le V^{\bar{u}}_{I-i}(G) - V^{\bar{u}}_{I-i}(G-y_i) && (7)\\ \to\ & p_i(y_i,(\hat{u}_{-i},\hat{b}_{-i})) \le p_i(y_i,(u_{-i},b_{-i})) && (8)\end{aligned}$$

where equality 6 is the definition of threshold price $p_i(y_i,(\hat{u}_{-i},\hat{b}_{-i}))$, and inequality 7 uses the fact that $V^{\bar{u}}_{I-i}(G) \ge V^{\bar{\hat{u}}}_{I-i}(G)$ (statement 3 of lemma 6) and $V^{\bar{\hat{u}}}_{I-i}(G-y_i) = V^{\bar{u}}_{I-i}(G-y_i)$ (statement 1 of lemma 6). Lastly, inequality 8 uses the definition of threshold price $p_i(y_i,(u_{-i},b_{-i}))$.

Consider the chosen allocation x^* at profile $(\hat{u},\hat{b})$. Note that by construction, j wins at x^*, so by step 3 of the allocation rule, truncated valuation maximization among $(I-i)$ from goods in $G-x^*_i$ involves bidder j, i.e., $x^{\bar{\hat{u}}}_{I-i,j}(G-x^*_i)=\bar{x}_j$. I have

$$\begin{aligned} V^{\bar{u}}_{I-i}(G) &\le V^{\bar{\hat{u}}}_{I-i}(G)+\delta && (9)\\ V^{\bar{u}}_{I-i}(G-x^*_i) &= V^{\bar{\hat{u}}}_{I-i}(G-x^*_i)+\delta && (10)\\ \to p_i(x^*_i,(u_{-i},b_{-i})) &\le p_i(x^*_i,(\hat{u}_{-i},\hat{b}_{-i})). && (11)\end{aligned}$$

In the above, inequality 9 comes from statement 3 of lemma 6, inequality 10 comes from statement 2 of lemma 6, noting that $x^{\bar{\hat{u}}}_{I-i,j}(G-x^*_i)=\bar{x}_j$. Lastly, inequality 11 is gotten by subtracting inequalities 9 and 10 side by side, and using the definition of threshold prices.

Finally,

$$\begin{aligned} u_i(x^*_i)-p_i(x^*_i,(\hat{u}_{-i},\hat{b}_{-i})) &> u_i(y_i)-p_i(y_i,(\hat{u}_{-i},\hat{b}_{-i})) && (12)\\ \to u_i(x^*_i)-p_i(x^*_i,(u_{-i},b_{-i})) &> u_i(y_i)-p_i(y_i,(\hat{u}_{-i},\hat{b}_{-i})) && (13)\\ \to u_i(x^*_i)-p_i(x^*_i,(u_{-i},b_{-i})) &> u_i(y_i)-p_i(y_i,(u_{-i},b_{-i})), && (14)\end{aligned}$$

where inequality 12 comes from the strict optimality of x^*_i for i at $(\hat{u},\hat{b})$, and inequalities 13 and 14 comes from inequalities 11 and 8 respectively.

By inequality 14, bundle y_i is not optimal for bidder i, a contradiction.

Therefore, $x^{\bar{u}}_{I-i,j}(G-y_i)=\bar{x}_j$, which means, by step 3 of the allocation rule, that j wins $\bar{x}_j$ at (u,b).

It is entirely symmetric to show that if j bids a truncated valuation below his threshold price, he loses. Alternatively, by the definition of threshold price, if j makes an unconstrained bid below his threshold price, he loses. Since the mechanism $\phi(\cdot)$ only cares about j's truncated valuation, such a bid is outcome-equivalent to a bid with truncated valuation below threshold price, so j must lose with such a bid. □

The variant of truncation VCG above works in the case of one multi-minded bidder because it explicitly sets out to satisfy the optimality condition for the multi-minded bidder through steps 1 and 2 of the allocation procedure. As shown in lemma 3, the optimality condition is necessary for incentive compatibility. If there are multiple multi-minded bidders, however, it may not be possible to satisfy the optimality condition for all such bidders simultaneously. In such a domain, there is no mechanism that is IR, GPO and GIC. Theorem 6 formalizes this idea with a counter example.

The following lemma will be useful in proving theorem 6. It states that a multi-minded bidder must be given his optimal bundle, given the threshold prices arising from Pareto optimal allocation rule. Similar to IC-typical reports, a report (u_{-i}, b_{-i}) is *PO-typical* if the set

$$_{po}\mathbb{U}_i = \{(u_i, b_i) : \phi(\cdot) \text{ is Pareto optimal at } ((u_i, b_i), (u_{-i}, b_{-i}))\}$$

contains almost all i's types.

Lemma 7. *Consider the domain in which a bidder i is multi-minded and constrained, and all other bidders are unconstrained and let $\phi(\cdot)$ be a mechanism that is IR, GPO and GIC in this domain. If report (u_{-i}, b_{-i}) is both IC-typical and PO-typical, then at any profile $(u, b) = ((u_i, b_i), (u_{-i}, b_{-i}))$ at which $\phi(\cdot)$ is incentive compatible, the allocation must be as described by the following procedure.*

1. *Let $\hat{p}_i(x_i, (u_{-i}, b_{-i})) = V^u_{I-i}(G) - V^u_{I-i}(G - x_i)$.*
2. *Allocation for bidder i is $\phi^a_i(u, b) = \arg\max_{x_i} v_i(x_i, \hat{p}_i(x_i, (u_{-i}, b_{-i})))$.*
3. *Allocation for other bidders is $\phi^a_{-i}(u, b) = x^u_{I-i}(G - x_i)$.*

Proof. Consider a report (u_{-i}, b_{-i}) that is both IC-typical and PO-typical, so there is a set $^{ic}_{po}\mathbb{U}_i$ containing almost all i's types, such that for any type (u_i, b_i) in $^{ic}_{po}\mathbb{U}_i$, $\phi(\cdot)$ is both incentive compatible and Pareto optimal at $((u_i, b_i), (u_{-i}, b_{-i}))$. In particular, there is also a set $_{po}\mathbb{U}^u_i$ containing almost all i's unconstrained types such that for any (u_i, b_i) in $_{p}\mathbb{U}_i$, $\phi(\cdot)$ is Pareto optimal at $((u_i, b_i), (u_{-i}, b_{-i}))$. At all profiles (u, b) in $_{po}\mathbb{U}^u_i \times (u_{-i}, b_{-i})$, Pareto optimality and perfect transferability of utility implies that the allocation must be $x^u_I(G)$, i.e., to maximize total valuation, at these profiles. It is then straightforward to establish that i's threshold price is given by $p_i(x_i, (u_{-i}, b_{-i})) = V^u_{I-i}(G) - V^u_{I-i}(G - x_i)$, which is the same as $\hat{p}_i(x_i, (u_{-i}, b_{-i}))$.

Now consider any profile $(u, b) = ((u_i, b_i), (u_{-i}, b_{-i}))$ at which $\phi(\cdot)$ is incentive compatible. Because (u_{-i}, b_{-i}) is IC-typical and $\phi(\cdot)$ is incentive compatible at (u, b), the optimality condition must hold (lemma 3), so step 2 of the procedure is established. Once i's allocation is determined, Pareto optimality and perfect transferability of utility among remaining unconstrained bidders dictate that the remaining goods must be allocated to maximize total valuation. Hence, step 3 holds. □

Now I can state and prove the following theorem, which describes the domain in which implementation of Pareto optimal outcome is not possible.

Theorem 6. *Consider the domain in which one bidder is multi-minded and constrained, another bidder is multi-minded and unconstrained, and all other bidders are unconstrained. There is no mechanism that is IR, GPO and GIC in this domain.*

Proof. I provide a sketch with the main ideas here and defer the formal proof to the appendix. Assume for now that the allocation procedure in lemma 7 applies to all profiles used in this proof sketch. Consider the profile in table 5,[12] where bidder 2 is the multi-minded constrained bidder and bidder 1 is an unconstrained multi-minded bidder. Using allocation procedure described in lemma 7, step 1 indicates threshold prices facing bidder 2 are 15, 10 and 25 for A, B, and AB respectively, and step 2 implies that bidder 2 wins B. A is then allocated to bidder 1 by step 3. The allocation is indicated by a box around the valuation.

I first argue that the same allocation procedure implies that the threshold prices for bidder 1 are 12 for A and 5 for B (as shown in table 5). Therefore, the allocation procedure does not satisfy the optimality condition, so the mechanism cannot be incentive compatible.

Table 5: A counter example

Bidder	Valuations			Budget	Threshold Prices		
	A	B	AB		A	B	AB
1	$\boxed{15}$	10	25	∞	12	5	
2	30	$\boxed{11}$	41	12	15	10	25
3	7	0	7	∞			

Table 6: Bidder 1's threshold prices

	Bidder 1's bids			Bidder 2's threshold prices		
	A	B	AB	A	B	AB
Bid 1	$\boxed{15}$	10	25	15	$\boxed{10}$	25
Bid 2	$\boxed{12 \pm \epsilon}$	0	$12 \pm \epsilon$	$12 \pm \epsilon$	$\boxed{0}$	$12 \pm \epsilon$
Bid 3	0	$\boxed{5 \pm \epsilon}$	$5 \pm \epsilon$	$\boxed{7}$	$5 \pm \epsilon$	$12 \pm \epsilon$

Claim 1: Bidder 1's threshold price for bundle A is 12. To see this, observe that if bidder 1 makes an unconstrained single-minded bid of $12 + \epsilon$ on A ("Bid 2" in table 6) for some small $\epsilon > 0$, he induces threshold prices facing bidder 2 to be $12 + \epsilon$ for both A and AB as shown in the table, so bidder 2 can only afford B and bidder 1 wins A consequently. If bidder 1 bids $12 - \epsilon$, however, then the resulting threshold prices facing bidder 2 imply that bidder 2 wins AB, so bidder 1 does not win A. Therefore, his threshold price for A is 12.

Claim 2: Bidder 1's threshold price for bundle B is 5. To see this, observe that if bidder 1 makes an unconstrained single-minded bid of $5 + \epsilon$ on B ("Bid 3" in table 6) for some small $\epsilon > 0$, he induces threshold prices facing bidder 2 to be as shown in the table, so by lemma 7 bidder 2 wins A and bidder 1 wins B consequently. If bidder 1 bids $5 - \epsilon$,

[12]The fact that valuations are additive at this profile is not essential to the proof.

however, then the resulting threshold prices facing bidder 2 imply that bidder 2 wins AB, so bidder 1 does not win B. Therefore, his threshold price for B is 5.

Therefore, $\phi(\cdot)$ does not satisfy the optimality condition for bidder 1 and so is not incentive compatible at the profile in table 5. $\square$

Theorem 6, in conjunction with other results in the current paper, describes the maximal domain in which Pareto optimal outcomes can be implemented in dominant strategy. A procedure is provided below to determine when implementation is possible, and if so, how.

- ◇ Is there any constrained bidder?
  - ◇ No → VCG is the unique mechanism.
  - ◇ Yes → Is there any multi-minded constrained bidder?
    - ◇ No → All constrained bidders ares single-minded. Truncation VCG is the unique mechanism (theorem 4).
    - ◇ Yes → There is at least one constrained multi-minded bidder. Is there another multi-minded bidder?
      - ◇ No → There is a mechanism (theorem 5).
      - ◇ Yes → There is no mechanism (theorem 6).

Essentially, the difficulty in the implementation resides with constrained bidders who are multi-minded. This is not to say that whenever there is a multi-minded constrained bidder such implementation is impossible. Theorem 5 shows that implementation is possible and presents a mechanism for such a domain. Nevertheless, one cannot step beyond the domain with one multi-minded constrained bidder. Theorem 6 implies that no mechanism exists when there are at least two multi-minded bidders, one of which is constrained. Technically, Pareto optimality implies certain threshold prices, and in the presence budget constraints there is a conflict among bidders in terms of their optimal bundles. In other words, the incentives of the bidders cannot be aligned through threshold prices induced by Pareto optimality.

6. Remarks on Generic Incentive Compatibility

The relaxation of *full* incentive compatibility to *generic* incentive compatibility is not always innocuous. For example, Agastya and Holden (2006) show that the set of equilibria changes markedly depending on whether incentive compatibility is required everywhere or almost everywhere. In the current paper, GIC is motivated by the observation on "knife-edge" profiles. Because full incentive compatibility seems to fail at apparently knife-edge profiles which constitute a set of measure zero, I posit the GIC condition which allows incentive compatibility to fail at a set of profiles of measure zero, and investigate mechanisms that satisfy GIC. My definition of GIC is agnostic as to which set of profiles incentive compatibility can be violated, because from the outset there is no obvious candidate set: a set of profiles can be knife-edge for one allocation rule, but generic for another allocation rule. Due to this indeterminacy, it could be that (1) the choice of such sets can affect the results, and, in particular, (2) GIC may be allowing too many incentive constraints to be violated, making it easier for candidate mechanisms to satisfy the three conditions (IR, GPO and GIC). I address these concerns below.

I start with the discussion on the positive results, and for ease of exposition, I restrict attention to the single-good setting. Though not explicit in the statements of theorem 1, the proof of the theorem shows that truncation VCG is incentive compatible on the set $\mathbb{U}_t$ of profiles where the maximal bidder is unique. In other words, as far as truncation VCG is concerned, GIC is invoked to relax incentive constraints at precisely only the set of profiles with multiple maximal bidders. In general, given a GIC mechanism, one can, in principle, derive exactly the set of profiles at which it is incentive compatible.

Consider now the task of characterizing GIC mechanisms. As mentioned above, *a priori* one does not know at what set of profiles the mechanism may violate incentive constraints, only that such a set has measure zero, so one cannot be explicit about this set. One could posit a specific set of profiles (such as $\mathbb{U}\backslash\mathbb{U}_t$), and look for the associated GIC mechanisms, but there is no clear justification for selecting such a set. However, *regardless* of what this set is, the mechanism must coincide with truncation VCG *a.e.* One interpretation is that even though GIC can be invoked to relax incentive constraints at any set of measure zero, such freedom does *not* give rise to an abundance of satisfactory mechanisms - in fact, all such mechanisms are essentially truncation VCG (modulo a set of measure zero).

For the negative result (theorem 6), it is straightforward to see that making GIC more or less demanding does not change the result. As long as incentive constraints are required to hold at almost all profiles, there is no satisfactory mechanism.

7. Conclusion

The current paper studies budgeted combinatorial auctions with a focus on incentive compatible mechanisms that implement Pareto optimal outcomes. I show that when the setting involves a single good or multiple goods but with single-minded bidders, truncation VCG is the unique mechanism that is individually rational, generically Pareto optimal and generically incentive compatible. As a result, truncation VCG can be used to attain Pareto optimality in single-good auctions, or multi-good auctions where the single-minded assumption is appropriate, such as auctioning of pollution rights, communication links in a tree or auto parts to buyers desiring a specific model (see Lehmann et al., 2002, and the references therein). Other potential applications are spectrum auctions where the auctioneer has sufficient information about a bidder's existing technology and wireless infrastructures to be confident that the bidder is only interested in only one specific spectrum bundle. It is worth noting, however, that, like VCG, truncation VCG has practical limitations, such as its sealed bid format, lack of revenue monotonicity and susceptibility to group manipulations (see Ausubel and Milgrom, 2006).

The paper also provides maximal domain results that describe when Pareto optimal outcomes can be implemented in dominant strategies. I show that the difficulty in implementation resides in the presence of multiple multi-minded bidders, one of whom is constrained: whenever such bidders are present, implementation is impossible. I also show that for domains in which such bidders are absent, implementation is possible via variants of truncation VCG. Collectively, these maximal domain results allow the auction designer to, given the setting at hand, determined whether implementation of Pareto optimal outcomes is possible and, if so, provide an implementing mechanism.

A. Appendix: Omitted proofs

A.1. Proof of lemma 2

A mechanism with optimality and threshold pricing essentially constructs personalized posted prices for each bidder, who gets his optimal bundle given these prices. Consider any profile $(u,b) \in U$, any bidder i, and suppose that bidder i's allocation at this profile is x_i. By threshold pricing, his payoff is $v_i(x_i, p_i(x_i,(u_{-i},b_{-i})))$. Consider any misreport by bidder i that results in allocation $\hat{x}_i$. Again, by threshold pricing, his payment is $p_i(\hat{x}_i,(u_{-i},b_{-i}))$, resulting in payoff of $v_i(\hat{x}_i, p_i(\hat{x}_i,(u_{-i},b_{-i})))$. By the optimality condition, however, $v_i(x_i, p_i(x_i,(u_{-i},b_{-i}))) \geq v_i(\hat{x}_i, p_i(\hat{x}_i,(u_{-i},b_{-i})))$. So this misreport does not benefit i. Therefore, i has no profitable deviation. □

A.2. Proof of lemma 5

Let (u,b) be given, and consider any such outcome (x,p) with $x = x_I^{\bar{u}}(G)$, and suppose in negation that another outcome (y,q) (potentially involving negative payments) Pareto dominates (x,p). If $y = x$, then Pareto dominance requires $q_i \leq p_i$ for all $i \in I$, leading to $\sum_{i\in I} q_i \leq \sum_{i\in I} p_i$. The left-hand-side is the seller's payoff at (y,q) which, by Pareto dominance, must be weakly greater than the right-hand-side, which is his payoff at (x,p). So equality must hold, implying that $q_i = p_i$ for all i, and consequently $(x,p) = (y,q)$, contradicting Pareto dominance. Therefore, $y \neq x$.

Let W be the set of bidders winning at both outcomes, W_x be the set of bidders winning at x only, and W_y be the set of bidders winning at y only. Let $p_C = \sum_{i\in C} p_i$ denote the total payment from bidders in a set C to the seller at (x,p). Similarly, q_C denotes the total payment from C at (y,q). Table 7 shows the distribution of payoffs and payments at the outcome (x,p). Note that bidders in W_y lose at (x,p) and pay exactly 0.

Table 7: Distribution of payoffs at outcome (x,p)

Bidders	W	W_x	W_y
Total Payoff	$V_W(x) - p_W$	$V_{W_x}(x) - p_{W_x}$	0
Total Payment	p_W	p_{W_x}	0

I first bound from above the total payment that bidders in W and W_y can make at (y,q). By Pareto dominance, bidders in W, who are winning the same bundles at x and y (by exactness), must not pay more at (y, q) than at (x, p). So $q_i \leq p_i$ for all $i \in W$, which implies

$$q_W \leq p_W. \tag{15}$$

Pareto dominance also requires that the outcome (y, q) is individually rational for bidders in W_y, so bidders in W_y can collectively pay a total

$$q_{W_y} \leq V^{\bar{u}}_{W_y}(y). \tag{16}$$

By the assumption of x as the unique truncated valuation maximizer and the fact that $x \neq y$, $V^{\bar{u}}_W(y) + V^{\bar{u}}_{W_x}(y) + V^{\bar{u}}_{W_y}(y) < V^{\bar{u}}_W(x) + V^{\bar{u}}_{W_x}(x) + V^{\bar{u}}_{W_y}(x)$. Noting that $V^{\bar{u}}_W(y) = V^{\bar{u}}_W(x)$ (by theness and exactness) and $V^{\bar{u}}_{W_x}(y) = V^{\bar{u}}_{W_y}(x) = 0$ yields $V^{\bar{u}}_{W_y}(y) < V^{\bar{u}}_{W_x}(x)$, which in turn combines with inequality 16 to give

$$q_{W_y} < V^{\bar{u}}_{W_x}(x). \tag{17}$$

Adding inequality 15 and inequality 17 yields

$$q_W + q_{W_y} < p_W + V^{\bar{u}}_{W_x}(x). \tag{18}$$

At (x, p), as can be seen from table 7, the total payoff to the seller and W_x is $p_W + V^u_{W_x}(x)$, which is weakly greater than $p_W + V^{\bar{u}}_{W_x}(x)$. By Pareto dominance, the outcome (y, q) must allocate a total valuation of at least $p_W + V^{\bar{u}}_{W_x}(x)$ to the seller and W_x. This total valuation must come in the form of payments from W and W_y, so these bidders must pay at least $p_W + V^{\bar{u}}_{W_x}(x)$, but this contradicts inequality 18.

Therefore, (y, q) cannot Pareto dominate (x, p), and (x, p) is Pareto optimal. □

A.3. Proof of theorem 3

The "sufficient" direction: In the single-minded domain, truncation VCG is IR, GPO and GIC.

Just like VCG in the unconstrained setting never makes a bidder pay more than his valuation, truncation VCG never makes a bidder pay more than his truncated valuation and is therefore individually rational. Recall the set $\mathbb{U}_t$ of profiles where the truncated valuation maximizer $x^{\bar{u}}_I(G)$ is unique. The complement of $\mathbb{U}_t$ consists of profiles where $x^{\bar{u}}_I(G)$ is multi-valued, and has measure zero,[13] so $\mathbb{U}_t$ contains almost all profiles. I show that truncation VCG is incentive compatible and Pareto optimal at profiles in $\mathbb{U}_t$.

Incentive compatibility on set $\mathbb{U}_t$ is shown first by verifying the conditions of threshold pricing and optimality in lemma 2. The allocation rule of truncation VCG induces threshold prices $p_i(x_i, (u_{-i}, b_{-i})) = V^{\bar{u}}_{I-i}(G) - V^{\bar{u}}_{I-i}(G - x_i)$, which coincides with the truncation

[13] Given a profile with multiple truncated valuation maximizers, there is some small perturbation of the valuations and budgets of the bidders which result in a profile whose truncated valuation maximizer is single-valued. Intuitively, the set of profiles with multiple truncated valuation maximizers have no "volume."

VCG payment that i makes if his allocation is x_i, so truncation VCG satisfies threshold pricing. To show optimality, it suffices to consider two possible allocations for i: (1) he wins his bundle of interest $\bar{x}_i$ or (2) he loses. At any profile (u,b) where the maximizer $x_I^{\bar{u}}(G)$ is unique, either (1) $\bar{u}_i(\bar{x}_i) + V_{I-i}^{\bar{u}}(G - \bar{x}_i) > V_{I-i}(G)$, which is equivalent to $\bar{u}_i(\bar{x}_i) - p_i(\bar{x}_i, (u_{-i}, b_{-i})) > 0$, or (2) $\bar{u}_i(\bar{x}_i) + V_{I-i}^{\bar{u}}(G - \bar{x}_i) < V_{I-i}(G)$, which is equivalent to $\bar{u}_i(\bar{x}_i) - p_i(\bar{x}_i, (u_{-i}, b_{-i})) < 0$. Truncation VCG allocates $\bar{x}_i$ to i in case (1) and nothing to i in case (2), and therefore satisfies optimality.

Pareto optimality on $\mathbb{U}_t$ follows from lemma 5, noting that truncation VCG is individually rational and, given truthful reporting, chooses the truncated valuation maximizer as allocation. □

The "necessary" direction: Any mechanism that is IR, GPO and GIC in the single-minded domain must coincide with truncation VCG a.e.

The proof parallels the proof sketch of theorem 2. Denote by $\mathbb{U}^k = \{(u,b) \in \mathbb{U} : |\{i : b_i < c_i\}| \leq k\}$ is the set of profiles containing at most k constrained bids, and recall that $\mathbb{U}_t$ is the set of profiles with unique truncated valuation maximizer. I induce on the number of constrained bids.

Claim 4 (Base case)**.** *At almost all profiles (u,b) in $\mathbb{U}^0$, $\phi^a(u,b) = x_I^{\bar{u}}(G)$.*

Proof. Consider the set $\mathbb{U}^0 \cap \mathbb{U}^t$, which contains almost all profiles in $\mathbb{U}^0$. At any profile (u,b) in this set, Pareto optimality and perfect transferability of utility (no bidder is constrained) implies that allocation must be $x_I^u(G)$ to maximize total surplus. Since no bidder is constrained, $x_I^u(G)$ coincides with $x_I^{\bar{u}}(G)$. □

Claim 5 (Inductive step - threshold prices)**.** *Suppose that at almost all profiles in $\mathbb{U}^k$, $\phi^a(u,b) = x_I^{\bar{u}}(G)$. Then at almost all profiles $(u,b) \in \mathbb{U}^{k+1}$, the threshold price for any bidder i making a constrained bid at the report (u_{-i}, b_{-i}) given by:*

$$p_i(x_i, (u_{-i}, b_{-i})) = V_{I-i}^{\bar{u}}(G) - V_{I-i}^{\bar{u}}(G - x_i).$$

Proof. Fix a bidder i, and let $\mathbb{U}_{-i}^k = \{(u_{-i}, b_{-i}) \in \mathbb{U}_{-i} : |\{j : j \neq i \text{ and } b_j < c_j\}| \leq k\}$ be the set of other bidders' reports containing at most k constrained bids. Recall that the report (u_{-i}, b_{-i}) is called IC-typical if the set ${}^{ic}\mathbb{U}_i = \{(u_i, b_i) : \phi(\cdot) \text{ is incentive compatible at } ((u_i, b_i), (u_{-i}, b_{-i}))\}$ contains almost all i's types. A report (u_{-i}, b_{-i}) in $\mathbb{U}_{-i}^k$ is called *tVCG-typical* if there is a set ${}_V\mathbb{U}_i^u$ containing almost all i's *unconstrained* types such that for any unconstrained type $(u_i, b_i) \in {}_V\mathbb{U}_i^u$, $\phi^a((u_i, b_i), (u_{-i}, b_{-i})) = x_I^{\bar{u}}(G)$. Let ${}_V^{ic}\mathbb{U}_{-i}^k$ be the set of reports in $\mathbb{U}_{-i}^k$ which are both tVCG-typical and IC-typical. Construct the set ${}_V^{ic}\mathbb{U}^{k+1}(i) = ({}_V^{ic}\mathbb{U}_{-i}^k \times \mathbb{U}_i) \cap {}^{ic}\mathbb{U}^{k+1}$ and ${}_V^{ic}\mathbb{U}^{k+1}(I) = \bigcap_{i \in I} {}_V^{ic}\mathbb{U}^{k+1}(i)$. By construction, any profile (u,b) in set ${}_V^{ic}\mathbb{U}^{k+1}(I)$ has three properties:

1. It contains at most $k+1$ constrained bids, so for any bidder i making a constrained bid, the report (u_{-i}, b_{-i}) contains at most k constrained bids and is tVCG-typical.

2. For any bidder i the report (u_{-i},b_{-i}) is IC-typical.
3. The mechanism $\phi(\cdot)$ is incentive compatible at (u,b).

Property (1) allows me to compute threshold prices (shown below), and properties (2) and (3) require the threshold pricing and optimality conditions in lemma 3 to hold.

I now argue that at any profile (u,b) in ${}^{c}_{V}\mathbb{U}^{k+1}(I)$, for any bidder i making a constrained bid the threshold price is:

$$p_i(x_i,(u_{-i},b_{-i})) = V^{\bar{u}}_{I-i}(G) - V^{\bar{u}}_{I-i}(G - x_i). \tag{19}$$

Because (u_{-i},b_{-i}) is both tVCG-typical and IC-typical, there is a set ${}^{ic}_{V}\mathbb{U}^{u}_{i}$ containing almost all i's unconstrained types such that for any unconstrained type $(u_i,b_i) \in {}^{ic}_{V}\mathbb{U}^{u}_{i}$, $\phi(\cdot)$ agrees with truncation VCG, allocation-wise, at the profile $((u_i,b_i),(u_{-i},b_{-i}))$ and is incentive compatible at this profile. Consider two cases: (1) $p_i(x_i,(u_{-i},b_{-i})) > V^{\bar{u}}_{I-i}(G) - V^{\bar{u}}_{I-i}(G - x_i)$ and (2) $p_i(x_i,(u_{-i},b_{-i})) < V^{\bar{u}}_{I-i}(G) - V^{\bar{u}}_{I-i}(G - x_i)$. For case (1), pick an unconstrained bid $(\hat{u}_i,\hat{b}_i)$ from ${}^{ic}_{V}\mathbb{U}^{u}_{i}$ such that $p_i(x_i,(u_{-i},b_{-i})) > \hat{u}_i(x_i) > V^{\bar{u}}_{I-i}(G) - V^{\bar{u}}_{I-i}(G - x_i)$.[14] By the choice of $(\hat{u}_i,\hat{b}_i)$, the allocation must be the truncation VCG allocation $x^{\bar{u}}_{I}(G)$ at $((\hat{u}_i,\hat{b}_i),(u_{-i},b_{-i}))$. Because $\hat{u}_i(x_i) > V^{\bar{u}}_{I-i}(G) - V^{\bar{u}}_{I-i}(G - x_i)$, $x^{\bar{u}}_{I}(G)$ allocates x_i to i. But this means that i wins x_i by bidding less than his threshold price, contradicting the definition of threshold prices.

For case (2), pick an unconstrained type $(\hat{u}_i,\hat{b}_i)$ in ${}^{ic}_{V}\mathbb{U}^{u}_{i}$ with $p_i(x_i,(u_{-i},b_{-i})) < \hat{u}_i(x_i) < V^{\bar{u}}_{I-i}(G) - V^{\bar{u}}_{I-i}(G - x_i)$.[15] By the assumed allocation rule, i loses at this profile, so his payoff is zero. However, by the choice of $(\hat{u}_i,\hat{b}_i)$, the mechanism $\phi(\cdot)$ must be incentive compatible at $((\hat{u}_i,\hat{b}_i),(u_{-i},b_{-i}))$. Hence, by lemma 1, i's payoff is bounded below by $v_i(x_i,p_i(x_i,(u_{-i},b_{-i}))) = \hat{u}_i(x_i) - p_i(x_i,(u_{-i},b_{-i})) > 0$, a contradiction.

To establish the claim, it remains to show that the set ${}^{ic}_{V}\mathbb{U}^{k+1}(I)$ contains almost all profiles in $\mathbb{U}^{k+1}$. The set of tVCG-typical reports in $\mathbb{U}^{k}_{-i}$ contains almost all reports in $\mathbb{U}^{k}_{-i}$.[16] The set of IC-typical reports in $\mathbb{U}^{k}_{-i}$ also contains almost all reports in $\mathbb{U}^{k}_{-i}$,[17] so the set of reports in $\mathbb{U}^{k}_{-i}$ which are both tVCG-typical and IC-typical, ${}^{ic}_{V}\mathbb{U}^{k}_{-i}$, contains almost all reports in $\mathbb{U}^{k}_{-i}$. The set ${}^{ic}_{V}\mathbb{U}^{k+1}(i) = ({}^{ic}_{V}\mathbb{U}^{k}_{-i} \times \mathbb{U}_i) \cap {}^{ic}\mathbb{U}^{k+1}$ then contains almost all profiles in $\mathbb{U}^{k+1}$, and the intersection of all such sets, ${}^{ic}_{V}\mathbb{U}^{k+1}(I)$, also contains almost all profiles. □

[14] This is always possible because ${}^{ic}_{V}\mathbb{U}^{u}_{i}$ contains almost all i's unconstrained types.

[15] Such a report exists because ${}^{ic}_{V}\mathbb{U}^{u}_{i}$ contains almost all unconstrained types.

[16] Suppose otherwise, then there must be a set of strictly positive measure of reports $\hat{U}^{k}_{-i}$ such that at any report $(u_{-i},b_{-i}) \in \hat{U}^{k}_{-i}$, there is a set of strictly positive measure of i's unconstrained type $\hat{U}^{u}_{i}$ such that for any $(u_i,b_i) \in \hat{U}_i$, the allocation is *not* $x^{\bar{u}}_{I}(G)$ at $((u_i,b_i),(u_{-i},b_{-i}))$. But this means there is a set of profiles in $\mathbb{U}^{k}$ with strictly positive measure at which the allocation is not $x^{\bar{u}}_{I}(G)$, contradicting the inductive hypothesis.

[17] See footnote 7.

Claim 6 (Inductive step - allocation)**.** *Suppose that at almost all profiles in $\mathbb{U}^k$, $\phi^a(u,b) = x_I^{\bar{u}}(G)$. Then at almost all profiles $(u,b) \in \mathbb{U}^{k+1}$, the allocation is $\phi^a(u,b) = x_I^{\bar{u}}(G)$.*

Proof. Now I show that at any profile (u,b) in ${}^{ic}_{V}\mathbb{U}^{k+1}(I)$ (defined above), the truncation VCG allocation $x^* = x_I^{\bar{u}}(G)$ must be chosen. Let C be the set of bidders making constrained bids at this profile. There are at most $k+1$ bidders in C. Let x^*_{I-C} denote the allocation for bidders in $I-C$, and $G - x^*_C$ denote the goods remaining for $I-C$, assuming that x^*_C is the allocation for bidders in C. Because x^* maximizes total truncated valuation, if x^*_C is allocated to C then x^*_{I-C} must attain the maximum truncated valuation possible from allocating $G - x^*_C$ to $I-C$. In other words, the principle of optimality implies that

$$x^*_{I-C} = x^{\bar{u}}_{I-C}(G - x^*_C). \tag{20}$$

For any bidder i making a constrained bid on a bundle of interest $\bar{x}_i$, if $x^*_i = \bar{x}_i \neq \emptyset$, then by the definition of x^* as the unique maximizer of truncated valuation, $\bar{u}_i(\bar{x}_i) > V^{\bar{u}}_{I-i}(G) - V^{\bar{u}}_{I-i}(G - \bar{x}_i)$. By equation 19, the right-hand-side of this inequality is his threshold price. In other words, i's truncated valuation exceeds his threshold price for his bundle of interest $\bar{x}_i$. By properties 2 and 3 of the set ${}^{c}_{V}\mathbb{U}^{k+1}(I)$ (see the proof of claim 5), the optimality condition in lemma 3 requires that he win $\bar{x}_i$. If $x^*_i = \emptyset$, then by the definition of x^*, $\bar{u}_i(\bar{x}_i) < V^{\bar{u}}_{I-i}(G) - V^{\bar{u}}_{I-i}(G - \bar{x}_i)$. So i's truncated valuation is less than his threshold price for his bundle, so he must lose. Hence, x^*_C describes the allocation for the bidders in C.

The remaining goods in $(G - x^*_C)$ must be allocated to bidders in $(I - C)$ (and potentially unallocated). Pareto optimality and perfect transferability of utility among bidders in $(I-C)$, who are all unconstrained, require that this allocation must maximize total valuation among bidders in $(I-C)$, i.e., $x^u_{I-C}(G - x^*_C)$ must be chosen. Because bidders in $I-C$ make unconstrained bids, $x^u_{I-C}(G - x^*_C)$ coincides with $x^{\bar{u}}_{I-C}(G - x^*_C)$ which, by equation 20, is simply x^*_{I-C}. □

Claims 4 through 6 establish that $\phi^p(\cdot)$ coincides with truncation VCG allocation *a.e.*. An argument similar to the proof of claim 5 establishes threshold prices at almost all profiles to be truncation VCG payments, and lemma 3 requires that these threshold prices are indeed payments for the mechanism $\phi(\cdot)$. This completes the proof. □

A.4. Proof of theorem 4

I first show that truncation VCG is IR, GPO and GIC. Individual rationality and generic incentive compatibility are straightforward to establish by verifying the threshold pricing and optimality conditions (which hold true by design of truncation VCG) on the set $\mathbb{U}_t$ of profiles whose truncated valuation maximizer is single-valued, and then invoking lemma 2.

I now show Pareto optimality on the set $\mathbb{U}_t$. Consider the truncation VCG outcome (x,p) with $x = x_I^{\bar{u}}(G)$, and suppose in negation that another outcome (y,q) (potentially

Table 8: Distribution of payoffs at (x,p)

Bidders	W	W_x	W_y	Z
Total Payoff	$V^u_W(x) - p_W$	$V^u_{W_x}(x) - p_{W_x}$	0	$V^u_Z(x) - p_Z$
Total Payment	p_W	p_{W_x}	0	p_Z

involving negative payments) Pareto dominates (x,p). Among the constrained bidders (who are all single-minded), let W be the set of bidders winning at both outcomes, W_x be the set of bidders winning at x only, and W_y be the set of bidders winning at y only. Let Z denote the unconstrained bidders. Let $p_C = \sum_{i\in C} p_i$ denote the total payment from bidders in a set C to the seller at (x,p). Similarly, q_C denotes the total payment from C at (y,q). Table 8 shows the distribution of the payoffs at the outcome (x,p). Note that bidders in W_y lose at (x,p) and pay exactly 0.

I first bound from above the total payment that bidders in W, W_y and Z can make. Similar to the proof of lemma 5, payment from W, q_W, is bounded by

$$q_W \le p_W, \tag{21}$$

and bidders in W_y can collectively pay a total of at most

$$q_{W_y} \le V^{\bar{u}}_{W_y}(y). \tag{22}$$

By Pareto dominance, the payment from Z, q_Z, needs to satisfy $V^u_Z(y) - q_Z \ge V^u_Z(x) - p_Z$, so

$$q_Z \le V^u_Z(y) - V^u_Z(x) + p_Z. \tag{23}$$

By the assumption of x as the unique truncated valuation maximizer, $V^{\bar{u}}_W(y) + V^{\bar{u}}_{W_x}(y) + V^{\bar{u}}_{W_y}(y) + V^{\bar{u}}_Z(y) < V^{\bar{u}}_W(x) + V^{\bar{u}}_{W_x}(x) + V^{\bar{u}}_{W_y}(x) + V^{\bar{u}}_Z(x)$. Noting that $V^{\bar{u}}_W(y) = V^{\bar{u}}_W(x)$ (by single-mindedness), $V^{\bar{u}}_{W_x}(y) = V^{\bar{u}}_{W_y}(x) = 0$ and that bidders in Z are unconstrained yields

$$V^u_Z(y) - V^u_Z(x) < V^{\bar{u}}_{W_x}(x) - V^{\bar{u}}_{W_y}(y). \tag{24}$$

Combining inequality 23 with inequality 24 yields

$$q_Z < V^{\bar{u}}_{W_x}(x) - V^{\bar{u}}_{W_y}(y) + p_Z. \tag{25}$$

Adding inequalities 21, 22 and 25 side by side gives

$$q_W + q_{W_y} + q_Z < p_W + V^{\bar{u}}_{W_x}(x) + p_Z. \tag{26}$$

At (x,p), as can be seen from table 8, the total payoff to the seller and W_x is $p_W + V^u_{W_x}(x) + p_Z$, which is weakly greater than $p_W + V^{\bar{u}}_{W_x}(x) + p_Z$. By Pareto dominance, the outcome (y,q) must allocate a total valuation of at least $p_W + V^{\bar{u}}_{W_x}(x) + p_Z$ to the seller and W_x. This total valuation comes in the form of payments from W, W_y and Z, so these bidders must pay at least $p_W + V^{\bar{u}}_{W_x}(x) + p_Z$, but this contradicts inequality 26.

Therefore, (y, q) cannot Pareto dominate (x, p), and truncation VCG is Pareto optimal on $\mathbb{U}_t$.

The proof that any mechanism that is IR, GPO and GIC must coincide with truncation VCG on generic profiles is identical to the induction proof of theorem 3. □

A.5. Proof of lemma 6

Throughout the proof, it is assumed that all truncated valuation maximizers are unique. This assumption is made only for notational convenience and is not necessary for the results to hold.

It will be useful to have the following simple observations, the proof of which is straightforward and omitted.

Claim 7. *Consider a set of bidders C and, an allocation x and a set of goods $\hat{G}$. The following statements are true.*

1. *If $x_j = \emptyset$, then $V^{\bar{u}}_C(x) = V^{\bar{\bar{u}}}_C(x)$.*
2. *If $j \notin C$, then $V^{\bar{u}}_C(x) = V^{\bar{\bar{u}}}_C(x)$ and $V^{\bar{u}}_C(\hat{G}) = V^{\bar{\bar{u}}}_C(\hat{G})$.*
3. *$V^{\bar{u}}_C(x) \geq V^{\bar{\bar{u}}}_C(x)$.*

Proof of statement 1 of lemma 6.

For any x in $X(\hat{G})$,

$$\begin{aligned} & V^{\bar{u}}_C(x^{\bar{u}}_C(\hat{G})) \geq V^{\bar{u}}_C(x) && (27)\\ \rightarrow\ & V^{\bar{\bar{u}}}_C(x^{\bar{u}}_C(\hat{G})) \geq V^{\bar{u}}_C(x) && (28)\\ \rightarrow\ & V^{\bar{\bar{u}}}_C(x^{\bar{u}}_C(\hat{G})) \geq V^{\bar{u}}_C(x) \geq V^{\bar{\bar{u}}}_C(x) && (29)\\ \rightarrow\ & V^{\bar{\bar{u}}}_C(x^{\bar{u}}_C(\hat{G})) \geq V^{\bar{\bar{u}}}_C(x), && (30)\end{aligned}$$

where inequality 27 is by optimality of $x^{\bar{u}}_C(\hat{G})$, inequality 28 uses the fact that $V^{\bar{u}}_C(x^{\bar{u}}_C(\hat{G})) = V^{\bar{\bar{u}}}_C(x^{\bar{u}}_C(\hat{G}))$ which is derived by statement 1 of claim 7, and the second inequality of 29 relies on statement 3 of claim 7.

Therefore, $x^{\bar{u}}_C(\hat{G}) = x^{\bar{\bar{u}}}_C(\hat{G})$. Since neither allocates any good to j, by statement 1 of claim 7, $V^{\bar{u}}_C(\hat{G}) = V^{\bar{\bar{u}}}_C(\hat{G})$. □

Proof of statement 2 of lemma 6.

Consider any allocation x in $X(\hat{G})$ and two cases: (1) $x_j = \bar{x}_j$ and (2) $x_j = \emptyset$. In case (1), note that

$$
\begin{aligned}
V_C^{\bar{u}}(x_C^{\bar{\bar{u}}}(\hat{G})) &= V_{C-j}^{\bar{u}}(x_{C,-j}^{\bar{\bar{u}}}(\hat{G})) + \bar{u}_j(\bar{x}_j) && (31)\\
&= V_{C-j}^{\bar{\bar{u}}}(x_{C,-j}^{\bar{\bar{u}}}(\hat{G})) + \bar{u}_j(\bar{x}_j) && (32)\\
&= V_{C-j}^{\bar{\bar{u}}}(\hat{G} - \bar{x}_j) + \bar{u}_j(\bar{x}_j) && (33)\\
&= V_{C-j}^{\bar{u}}(\hat{G} - \bar{x}_j) + \bar{u}_j(\bar{x}_j) && (34)\\
&\geq V_{C-j}^{\bar{u}}(x_{-j}) + \bar{u}_j(\bar{x}_j) = V_C^{\bar{u}}(x). && (35)
\end{aligned}
$$

In the above, equality 31 simply decomposes $x_C^{\bar{\bar{u}}}(\hat{G})$ to allocation for bidders other than j, $x_{C,-j}^{\bar{\bar{u}}}(\hat{G})$, and allocation to j, $\bar{x}_j$. Equality 32 relies on statement 2 of claim 7. Equality 33 uses the fact that, by principle of optimality, $x_{C,-j}^{\bar{\bar{u}}}(\hat{G})$ maximizes truncated valuation among bidders in $C - j$ from goods in $\hat{G} - \hat{x}_j$. Equality 34 relies on statement 2 of claim 7. Inequality 35 use the optimality of $V_{C-j}^{\bar{u}}(\hat{G} - \bar{x}_j)$ and the decomposition of x into x_{-j} and $\bar{x}_j$. Hence, $V_C^{\bar{u}}(x_C^{\bar{\bar{u}}}(\hat{G})) \geq V_C^{\bar{u}}(x)$.

In case (2), note that

$$
\begin{aligned}
& V_C^{\bar{u}}(x_C^{\bar{\bar{u}}}(\hat{G})) \geq V_C^{\bar{\bar{u}}}(x_C^{\bar{\bar{u}}}(\hat{G})) && (36)\\
\rightarrow\ & V_C^{\bar{u}}(x_C^{\bar{\bar{u}}}(\hat{G})) \geq V_C^{\bar{\bar{u}}}(x_C^{\bar{\bar{u}}}(\hat{G})) \geq V_C^{\bar{\bar{u}}}(x) && (37)\\
\rightarrow\ & V_C^{\bar{u}}(x_C^{\bar{\bar{u}}}(\hat{G})) \geq V_C^{\bar{\bar{u}}}(x) = V_C^{\bar{u}}(x) && (38)
\end{aligned}
$$

In the above, inequality 36 relies on statement 3 of claim 7. Inequality 37 uses the optimality of $x_C^{\bar{\bar{u}}}(\hat{G})$, and inequality 38 statement 1 of claim 7. Hence, $V_C^{\bar{u}}(x_C^{\bar{\bar{u}}}(\hat{G})) \geq V_C^{\bar{u}}(x)$.

So in both cases, $V_C^{\bar{u}}(x_C^{\bar{\bar{u}}}(\hat{G})) \geq V_C^{\bar{u}}(x)$, so $x_C^{\bar{u}}(\hat{G}) = x_C^{\bar{\bar{u}}}(\hat{G})$. Using the fact that $x_{C,j}^{\bar{\bar{u}}}(\hat{G}) = \bar{x}_j$ and $\bar{u}_j(\bar{x}_j) = \bar{\bar{u}}_j(\bar{x}_j) + \delta$, it is easy to see that $V_C^{\bar{u}}(\hat{G}) = V_C^{\bar{\bar{u}}}(\hat{G}) + \delta$. □

Proof of statement 3 of lemma 6.

To show the first part of the inequality, note that for any x in $X(\hat{G})$ with $x_j = \bar{x}_j$,

$$
\begin{aligned}
V_C^{\bar{\bar{u}}}(x) + \delta &= V_{C-j}^{\bar{\bar{u}}}(x_{-j}) + \bar{\bar{u}}_j(\bar{x}_j) + \delta && (39)\\
&= V_{C-j}^{\bar{u}}(x_{-j}) + \bar{u}_j(\bar{x}_j) && (40)\\
&= V_C^{\bar{u}}(x) && (41)
\end{aligned}
$$

where equality 39 simply decomposes x to x_{-j} and $\bar{x}_j$, equality 40 relies on statement 2 of claim 7, and the assumption that $\bar{u}_j(\bar{x}_j) = \bar{\bar{u}}_j(\bar{x}_j) + \delta$.

Note also that for any x in $X(\hat{G})$ with $x_j = \emptyset$, by statement 1 of claim 7, $V_C^{\bar{\bar{u}}}(x) = V_C^{\bar{u}}(x)$, so

$$V_C^{\bar{\bar{u}}}(x) + \delta = V_C^{\bar{u}}(x) + \delta > V_C^{\bar{u}}(x) \tag{42}$$

By equations 41 and 42, $V_C^{\bar{\bar{u}}}(x) + \delta \geq V_C^{\bar{u}}(x)$ for any x, so taking the maximum over all x in $X(\hat{G})$ gives $V_C^{\bar{\bar{u}}}(\hat{G}) + \delta \geq V_C^{\bar{u}}(\hat{G})$.

Similarly, by statement 3 of claim 7, $V_C^{\bar{u}}(x) \geq V_C^{\bar{\bar{u}}}(x)$ for any x, so taking the maximum over all x gives $V_C^{\bar{u}}(\hat{G}) \geq V_C^{\bar{\bar{u}}}(\hat{G})$. □

A.6. Proof of generic Pareto optimality for theorem 5

Table 9: Distribution of payoffs at outcome (x, p)

Bidders	i	W	W_x	W_y
Total Payoff	$u_i(x_i) - p_i$	$V_W^u(x) - p_W$	$V_{W_x}^u(x) - p_{W_x}$	0
Total Payment	p_i	p_W	p_{W_x}	0

Let (u, b) be a profile in $\mathbb{U}_\phi$ and consider the outcome (x, p) chosen by the mechanism $\phi(\cdot)$ at this profile, and suppose in negation that another outcome (y, q) Pareto dominates (x, p). Among the single-minded bidders, let W be the set of bidders winning at both outcomes, W_x be the set of bidders winning at x only, and W_y be the set of bidders winning at y only. Let $p_C = \sum_{i \in C} p_i$ be the total payment from bidders in C to the seller at (x, p). Similarly, q_C denotes the total payment from C at (y, q). Table 9 shows how valuations and payments are distributed at the outcome (x, p). Note that bidders in W_y lose and pay exactly 0.

I first bound from above the total payment that bidders in W, W_y and i can make at (y, q). Similar to the proof of lemma 5, payment from W, q_W, is bounded by

$$q_W \leq p_W, \tag{43}$$

and bidders in W_y can collectively pay a total of at most

$$q_{W_y} \leq V_{W_y}^{\bar{u}}(y). \tag{44}$$

Consider bidder i. By step 2 of the allocation procedure and the fact that x_i is chosen , I have

$$u_i(x_i) - p_i(x_i, (u_{-i}, b_{-i})) > u_i(y_i) - p_i(y_i, (u_{-i}, b_{-i})).$$

Substituting in the expression for threshold prices yields

$$u_i(x_i) - V_{I-i}^{\bar{u}}(G) + V_{I-i}^{\bar{u}}(G - x_i) > u_i(y_i) - V_{I-i}^{\bar{u}}(G) + V_{I-i}^{\bar{u}}(G - y_i).$$

The common term $V_{I-i}^{\bar{u}}(G)$ can be removed. By assumption of x as truncated valuation maximizing, $V_{I-i}^{\bar{u}}(G - x_i) = V_W^{\bar{u}}(x) + V_{W_x}^{\bar{u}}(x)$. By definition, $V_{I-i}^{\bar{u}}(G - y_i) \geq V_W^{\bar{u}}(y) +$

$V^{\bar{u}}_{W_y}(y)$. So the above inequality leads to

$$u_i(x_i) + V^{\bar{u}}_W(x) + V^{\bar{u}}_{W_x}(x) > u_i(y_i) + V^{\bar{u}}_W(y) + V^{\bar{u}}_{W_y}(y).$$

Because $V^{\bar{u}}_W(x) = V^{\bar{u}}_W(y)$ by single-mindedness, the above inequality simplifies to

$$V^{\bar{u}}_{W_y}(y) < u_i(x_i) + V^{\bar{u}}_{W_x}(x) - u_i(y_i),$$

which combines with inequality 44 to yield

$$q_{W_y} < u_i(x_i) + V^{\bar{u}}_{W_x}(x) - u_i(y_i). \tag{45}$$

To guarantee that i is not worse off at (y, q), it must be that $u_i(y_i) - q_i \geq u_i(x_i) - p_i$, which means

$$q_i \leq u_i(y_i) - u_i(x_i) + p_i. \tag{46}$$

Adding up inequalities 43, 45, and 46 yields

$$q_W + q_{W_y} + q_i < p_W + V^{\bar{u}}_{W_x}(x) + p_i. \tag{47}$$

At (x, p), as can be seen from table 9, the total payoff to the seller and W_x is $p_W + V^{u}_{W_x}(x) + p_i$, which is weakly greater than $p_W + V^{\bar{u}}_{W_x}(x) + p_i$. By Pareto dominance, the outcome (y, q) must allocate a total valuation of at least $p_W + V^{\bar{u}}_{W_x}(x) + p_i$ to the seller and W_x. This total valuation comes in the form of payments from W, W_y and i, so these bidders must pay at least $p_W + V^{\bar{u}}_{W_x}(x) + p_i$, but this contradicts inequality 47.

Therefore, (y, q) cannot Pareto dominate (x, p). □

A.7. Proof of theorem 6

Suppose to the contrary that there is a mechanism $\phi(\cdot)$ that is IR, GPO and GIC. Consider a specific domain with 3 bidders, bidder 1 being multi-minded and unconstrained, bidder 3 being single-minded and unconstrained, and bidder 2 being multi-minded and constrained. Let ${}^{ic}_{po}\mathbb{U}$ be the set of profiles at which the mechanism is both Pareto optimal and incentive compatible. Fix a bidder i, and consider the set of reports ${}^{ic}_{po}\mathbb{U}_{-i}$ which are both IC-typical and PO-typical. Construct the sets of profiles ${}^{ic}_{po}\mathbb{U}(i) = ({}^{ic}_{po}\mathbb{U}_{-i} \times \mathbb{U}_i) \cap {}^{ic}_{po}\mathbb{U}$ and ${}^{ic}_{po}\mathbb{U}(I) = \bigcap_{i \in I} {}^{ic}_{po}\mathbb{U}(i)$. By construction, at any profile (u, b) in ${}^{ic}_{po}\mathbb{U}(I)$, $\phi(\cdot)$ is incentive compatible and, for any bidder i, the report (u_{-i}, b_{-i}) is both IC-typical and PO-typical. This implies that at all profiles in ${}^{ic}_{po}\mathbb{U}(I)$, the allocation rule described in lemma 7 and the optimality condition in lemma 3 must hold. It is easy to show that ${}^{ic}_{po}\mathbb{U}_{-i}$ contains almost all reports (using argument similar to footnote 7). Hence, ${}^{ic}_{po}\mathbb{U}(i)$ contains almost all profiles and so does ${}^{ic}_{po}\mathbb{U}(I)$.

Some additional set manipulation is needed. A report (u_{-i}, b_{-i}) is ${}^{ic}_{po}\mathbb{U}(I)$-*typical* if the set ${}^*\mathbb{U}_i = \{(u_i,b_i) : ((u_i,b_i),(u_{-i},b_{-i})) \in {}^{ic}_{po}\mathbb{U}(I)\}$ contains almost all i's types. Let ${}^*\mathbb{U}_{-i}$ be the set of ${}^{ic}_{po}\mathbb{U}(I)$-typical reports. Similar to above, construct the set ${}^*\mathbb{U}(i) = ({}^*\mathbb{U}_{-i} \times \mathbb{U}_i) \cap {}^{ic}_{po}\mathbb{U}(I)$ and ${}^*\mathbb{U}(I) = \bigcap_{i \in I} {}^*\mathbb{U}(i)$. Since almost all reports are ${}^{ic}_{po}\mathbb{U}(I)$-typical,[18] the set ${}^*\mathbb{U}(I)$ contains almost all profiles.

Suppose for now that the profile in table 5 is in ${}^*\mathbb{U}(I)$. Because the report $(u,b)_{-1}$ is ${}^{ic}_{po}\mathbb{U}(I)$-typical, there is a set ${}^*\mathbb{U}_1$ containing almost all bidder 1's types such that for any type $(u,b)_1$ in ${}^*\mathbb{U}_1$, the profile $((u,b)_1,(u,b)_{-1})$ is in ${}^{ic}_{po}\mathbb{U}(I)$, which means that the allocation must be as described in lemma 7.

The threshold prices facing bidder 2 are as listed in table 5. Given these threshold prices, bidder 2 wins B, his optimal bundle. The remaining good A must be allocated to bidder 1 to maximize valuation. I claim that the allocation rule described in lemma 7 implies that the threshold prices for bidder 1 are 12 for A and 5 for B as shown in table 5 and therefore, this allocation rule does not satisfy the optimality condition in lemma 3 and consequently $\phi(\cdot)$ is not incentive compatible at this profile.

Table 10: Bidder 1's threshold prices

	Bidder 1's bids			Bidder 2's threshold prices		
	A	B	AB	A	B	AB
Bid 1	15	10	25	15	$\boxed{10}$	25
Bid 2	$\boxed{12+\epsilon}$	δ	$12+\epsilon+\delta+\eta$	$12+\epsilon+\eta$	$\boxed{\delta+\eta}$	$12+\epsilon+\delta+\eta$
Bid 3	$12-\epsilon+\delta$	η	$12-\epsilon+\delta+\eta+\nu$	$12-\epsilon+\delta+\nu$	$\eta+\nu$	$\boxed{12-\epsilon+\delta+\eta+\nu}$
Bid 4	δ	$\boxed{5+\epsilon}$	$5+\eta$	$\boxed{7}$	$5+\epsilon$	$12+\epsilon$
Bid 5	δ	$5-\epsilon+\eta$	$5-\epsilon+\eta+\nu$	7	$5-\epsilon+\eta$	$\boxed{12-\epsilon+\eta}$

Claim 1: Bidder 1's threshold price for bundle A, $p_1(A)$, is 12.

Proof of claim 1. Pick a type described by "Bid 2" in table 10 from ${}^*\mathbb{U}_1$ for some small $\epsilon,\delta,\eta > 0$. This is possible because ${}^*\mathbb{U}_1$ contains almost all profiles. By construction, the resulting profile is in ${}^{ic}_{po}\mathbb{U}(I)$, so the allocation rule described in lemma 7 must hold. At this bid, the threshold prices facing bidder 2 are computed and shown on the same table. Note that only B is affordable for bidder 2. According to the allocation procedure in lemma 7, bidder 2 wins B and bidder 1 wins A. So bidder 1's threshold price for A is at most 12.

Suppose that his threshold for A is strictly less than 12, so $p_1(A) = 12 - \epsilon$ for some $\epsilon > 0$. Similar to the above approach, pick a type from ${}^*\mathbb{U}_1$ described by "Bid 3" with a choice of small $\delta,\eta,\nu > 0$ such that $12 - \epsilon + \delta + \eta + \nu < 12$. At this bid, the resulting threshold prices for bidder 2 and the optimality condition mean that bidder 2 wins AB,[19] so bidder 1 does not win anything and gets zero payoff. However, this contradicts lemma

[18]This can be shown using the logic similar to footnote 7, using the fact that ${}^{ic}_{po}\mathbb{U}(I)$ contains almost all profiles.

[19]This is because, by choice of δ,η,ν, the threshold price for AB is strictly less than 12, so AB is the optimal bundle for bidder 2, given his valuations.

1, which states that his payoff must be at least the payoff from getting A at the threshold price $12 - \epsilon$, which is equal to $\delta > 0$.

Therefore, bidder 1's threshold price for bundle A is exactly 12, establishing claim 1.

Claim 2: Bidder 1's threshold price for bundle B, $p_1(B)$, is 5.

*Proof of claim 2.*The proof parallels the above. Pick a type described by "Bid 4" in table 10 from $^*\mathbb{U}_1$ for some small $\epsilon, \delta, \eta > 0$. At this bid, the threshold prices facing bidder 2 are computed and shown on the same table. Given these prices, bidder 2 wins A and bidder 1 wins B. So bidder 1's threshold price for B is at most 5.

Suppose that his threshold for B is strictly less than 5, so $p_1(B) = 5 - \epsilon$ for some $\epsilon > 0$. Pick a type from $^*\mathbb{U}_1$ described by "Bid 5" with a choice of small $\delta, \eta, \nu > 0$ such that $\epsilon > \eta$. At this bid, the resulting threshold prices for bidder 2 and the optimality condition mean that bidder 2 wins AB,[20] so bidder 1 does not win anything and gets zero payoff. However, this contradicts lemma 1, which states that his payoff must be at least the payoff from getting B at the threshold price $5 - \epsilon$, which is equal to $\eta > 0$.

Therefore, bidder 1's threshold price for B is 5, establishing claim 2.

I have assumed that the profile in table 5 is in the set $^*\mathbb{U}(I)$, while this may not necessarily be the case. However, because $^*\mathbb{U}(I)$ contains almost all profiles, it is always possible to pick a profile that allows the proof to work by, for example, "wiggling" the chosen profile a bit.

Therefore, $\phi(\cdot)$ is not generically incentive compatible. □

References

Agastya, M. and Holden, R. (2006). Incentive compatibility: Everywhere vs. almost everywhere incentive compatibility. Available at SSRN: https://ssrn.com/abstract=927403.

Ausubel, L. M. (2004). An efficient ascending-bid auction for multiple objects. *American Economic Review*, 94(5):1452–1475.

Ausubel, L. M. and Milgrom, P. (2006). The lovely but lonely Vickrey auction. *Combinatorial auctions*, 17:22–26.

Bhattacharya, S., Conitzer, V., Munagala, K., and Xia, L. (2010). Incentive compatible budget elicitation in multi-unit auctions. In *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*, pages 554–572. Society for Industrial and Applied Mathematics.

Borgs, C., Chayes, J., Immorlica, N., Mahdian, M., and Saberi, A. (2005). Multi-unit auctions with budget-constrained bidders. In *Proceedings of the 6th ACM conference on Electronic commerce*, pages 44–51. ACM.

Bulow, J., Levin, J., and Milgrom, P. (2009). Winning play in spectrum auctions. Technical report, National Bureau of Economic Research.

Che, Y.-K. and Gale, I. (1996). Expected revenue of all-pay auctions and first-price sealed-bid auctions with budget constraints. *Economics Letters*, 50(3):373–379.

Che, Y.-K. and Gale, I. (2000). The optimal mechanism for selling to a budget-constrained buyer. *Journal of Economic Theory*, 92(2):198–233.

Dobzinski, S., Lavi, R., and Nisan, N. (2012). Multi-unit auctions with budget limits. *Games and Economic Behavior*, 74(2):486–503.

[20]This is because, by choice of η, the threshold price for AB is strictly less than 12, so AB is the optimal bundle for bidder 2, given his valuations.

Fiat, A., Leonardi, S., Saia, J., and Sankowski, P. (2011). Single valued combinatorial auctions with budgets. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 223–232. ACM.

Goel, G., Mirrokni, V., and Paes Leme, R. (2012). Polyhedral clinching auctions and the adwords polytope. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 107–122. ACM.

Gul, F. and Stacchetti, E. (2000). The English auction with differentiated commodities. *Journal of Economic theory*, 92(1):66–95.

Hafalir, I. E., Ravi, R., and Sayedi, A. (2012). A near Pareto optimal auction with budget constraints. *Games and Economic Behavior*, 74(2):699–708.

Laffont, J.-J. and Robert, J. (1996). Optimal auction with financially constrained buyers. *Economics Letters*, 52(2):181–186.

Lavi, R. and May, M. (2012). A note on the incompatibility of strategy-proofness and pareto-optimality in quasi-linear settings with public budgets. *Economics Letters*, 115(1):100–103.

Lehmann, D., Oćallaghan, L. I., and Shoham, Y. (2002). Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM (JACM)*, 49(5):577–602.

Milgrom, P. (2000). Putting auction theory to work: The simultaneous ascending auction. *Journal of Political Economy*, 108(2):245–272.

Milgrom, P. and Segal, I. (2014). Deferred-acceptance auctions and radio spectrum reallocation. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 185–186. ACM.

Pai, M. M. and Vohra, R. (2014). Optimal auctions with financially constrained buyers. *Journal of Economic Theory*, 150:383–425.

Wilson, R. B. (1993). *Nonlinear pricing*. Oxford University Press on Demand.

Yokoo, M. (2003). Characterization of strategy/false-name proof combinatorial auction protocols: Price-Oriented, Rationing-Free protocol. In *IJCAI*, pages 733–742.