

Fair Payments for Efficient Allocations in Public Sector Combinatorial Auctions

Robert W. Day

Operations and Information Management, School of Business, University of Connecticut,
Storrs, Connecticut 06269, bob.day@business.uconn.edu

S. Raghavan

Decision and Information Technologies, The Robert H. Smith School of Business, University of Maryland,
College Park, Maryland 20742, raghavan@umd.edu

Motivated by the increasing use of auctions by government agencies, we consider the problem of fairly pricing public goods in a combinatorial auction. A well-known problem with the incentive-compatible Vickrey-Clarke-Groves (VCG) auction mechanism is that the resulting prices may not be in the core. Loosely speaking, this means the payments of the winners could be so low, that there are bidders who would have been willing to pay more than the payments of the winning bidders. Clearly, this “unfair” outcome is unacceptable for a public sector auction. Recent advances in auction theory suggest that combinatorial auctions resulting in efficient outcomes and bidder-Pareto-optimal core payments offer a viable practical alternative to address this problem.

This paper confronts two critical issues facing the bidder-Pareto-optimal core payment. First, motivated to minimize a bidder’s ability to benefit through strategic manipulation (through collusive agreement or unilateral action), we demonstrate the strength of a mechanism that minimizes total payments among all such auction outcomes, narrowing the previously broad solution concept. Second, we address the computational difficulties of achieving these outcomes with a constraint-generation approach, promising to broaden the range of applications for which bidder-Pareto-optimal core pricing achieves a comfortably rapid solution.

Key words: combinatorial auctions; core allocations; bidder-Pareto optimality; constraint generation; VCG payments; proxy auctions

History: Accepted by Linda V. Green, public sector applications; received February 1, 2005. This paper was with the authors 5½ months for 2 revisions. Published online in *Articles in Advance* July 20, 2007.

1. Introduction

Classic auction theory shows that if a single item is auctioned via the submission of sealed-bid price offers, then bidders can be expected to report their bids honestly under a second-price mechanism (i.e., one in which the highest bid wins and pays the amount of the second-highest bid). Such a mechanism is said to be *incentive compatible*, meaning that truth-telling is a dominant strategy for every player, and that there is no incentive for unilateral deviation from the truth-telling strategy.

The analogous second-price mechanism for combinatorial auctions (in which bidders bid on combinations of items) is the well-known Vickrey-Clarke-Groves (VCG) mechanism, which has also been shown to be incentive compatible (for primary sources see Clarke 1971, Groves 1973, and Vickrey 1961). In the general version of this auction mechanism, bidders each submit a price for each possible combination of items. Winners are chosen to maximize the combined social value of awarded bundles (with no item going to more than one bidder); such

an allocation is referred to as *efficient*. Each winner in the VCG auction then pays an amount less than her bid on the bundle she is awarded, receiving a discount from her actual bid that is calculated to eliminate her ability to gain from falsifying her preferences. In particular, each winning bidder receives a discount equal to the difference in value between the efficient solution with all bidders and the efficient solution in the absence of that particular bidder. The problem of finding an efficient allocation for a given set of bids is known as the *winner-determination problem*, and is often computationally difficult (i.e., \mathcal{NP} -hard; see Rothkopf et al. 1998)

Although the VCG mechanism is widely discussed in the auction literature, its drawbacks are so numerous that it is rarely, if ever, used in practice. Indeed, there are too many problems to name here; we therefore direct the reader to a few good references on the problems with VCG mechanisms (see Ausubel and Milgrom 2002, Rothkopf and Harstad 1995, Rothkopf et al. 1990, and Sakurai et al. 2000). Principle among these drawbacks is that the VCG payments may be

so low (in a forward auction) that the outcome is not a “core” allocation (defined formally in the next section). Roughly stated, this means that a dissatisfied coalition of bidders may be able to suggest an auction outcome that is preferred by all the members of the coalition *and* the seller. This situation is clearly unsatisfying, especially when the seller is a government agency assigned the duty of fair allocation, and the winners have not paid enough to establish themselves as the fair recipients of the auction items relative to the competition.

Given that the VCG mechanism is undesirable in practice, one may next ask: What other payment mechanism can be used in a combinatorial auction? The most obvious alternative is a first-price (or pay-as-bid) mechanism, but this too has drawbacks. In a sealed-bid combinatorial auction, this payment rule encourages the bidders to submit bids that are just barely enough to achieve the efficient allocation. Indeed, bidding higher than the minimum amount necessary to secure a particular bundle simply gives more money to the seller in the pay-as-bid scenario. But if each player is trying to predict the minimum amount needed to win his efficiently awarded bundle, uncertainty about the bids of others makes this a precarious endeavor. With incomplete information about the preferences of others, trying to bid the minimum possible amount will often lead to bidding not quite enough, leading the auction mechanism to miss the efficient outcome as a consequence.

Unsatisfied with the two “extreme” possibilities for a payment mechanism (first-price and VCG), is there some other “second-price” combinatorial auction payment mechanism that maintains most of the desirable properties of the second-price mechanism for a single-item auction, while not suffering from the undesirable properties of the VCG mechanism? Stated differently, we would like a mechanism that encourages truthful bidding by charging less than the full amount bid on a bundle, but does not suffer from the extremely low seller revenues of the VCG mechanism.

Few combinatorial auction mechanisms other than VCG have been studied extensively that result in “second-prices,” payments that are typically lower than what was bid and are determined by the bids of others. One prominent example was proposed recently by Ausubel and Milgrom (2002),¹ an *ascending proxy auction* providing a viable practical alternative to the virtually unusable VCG mechanism for real-world combinatorial auction applications. Assuring competitive pricing and economic efficiency, the

proxy auction terminates at a desirable “bidder-Pareto-optimal core outcome” (defined formally in the next section). This presents an especially attractive solution for several public-good allocation problems, where the need for a “socially acceptable” outcome outweighs revenue maximization for the seller.

Designed as a proxy implementation of an iterative combinatorial auction, bidders in this mechanism are insulated against the dangers of bidding more than the minimum amount necessary to win a particular bundle, as the proxy submits bid information only at minimum increments. The particular algorithm for winner/payment determination in the ascending proxy auction (as first proposed by Ausubel and Milgrom 2002) is interpreted by de Vries et al. (2007) as an implementation of the subgradient algorithm: at each moment in the process, the algorithm solves an \mathcal{NP} -hard set-packing problem, determining a nonoverlapping set of bundles from all those demanded by the bidders at some current set of prices, and a net utility maximizing bundle for each bidder. Each nonanonymous bundle price is then adjusted up by one increment for each bundle that is not contained in the “seller’s choice” allocation. This process corresponds to incremental stepping in the direction of a subgradient. As is commonly the case with the subgradient algorithm, convergence is slow and depends critically on the choices of step size.

Compounding the problem of rapid computation is that each iteration of this algorithm requires the solution of an integer program (IP) for winner determination. Although each IP can be solved reasonably quickly given advanced computational techniques (see, for example, Günlük et al. 2005) and software (like CPLEX and XPRESS, for example), the repeated solution and slow convergence limits practical implementation (see Hoffman et al. 2006).

Even since the relatively new development of the ascending proxy auction, a few methods for arriving at the same types of outcomes more rapidly have been proposed, including the technique developed in this paper. Hoffman et al. (2006) give experimental evidence for slow performance of the subgradient-type ascending proxy implementation and show how to improve the computational speed of the proxy auction’s iterative implementation by starting at the VCG solution, and by using a more sophisticated adjustment (scaling) of the price increment. Their technique requires the switch to a sealed-bid (nonproxy) auction allowing the auctioneer to use all of the information reported by the bidders to solve instances of the winner determination problem. Given that the auction mechanism can be trusted not to make the bid information public (which may be legally enforced or entrusted to a neutral third party), we argue that such revelation is reasonable and should be implemented

¹ A closely related mechanism is described by Parkes (2001), while Hoffman et al. (2006) and Wurman et al. (2004) investigate variations on this theme.

for large-scale applications because access to the full set of bid reports may greatly accelerate computational performance. Wurman et al. (2004) provide a different approach to accelerating the proxy implementation by computing the “inflection” or “change” points in the iterative auction’s price trajectory. Their technique also requires the full release of bid information to the auctioneer, rather than withholding it in a proxy. In §5, we look at these methods in greater depth and compare them to the technique presented here.

More recently, Ausubel et al. (2006) provided a practical *clock-proxy* auction design that terminates with bidders submitting combinatorial bids in a *proxy* auction after some initial rounds designed to reveal price information. The final proxy phase is strategically equivalent to a sealed-bid combinatorial auction, justifying our focus on sealed-bid price mechanisms in the current treatment. Given the widespread applicability of their technique in public sector markets for electricity generation, spectrum licenses, and airport landing-slot rights (to name a few), the ability to produce final payments more rapidly and to better understand the selection of payment outcomes among the various possible core solutions may have a deep impact on our ability to successfully implement these market mechanisms. Indeed, researchers associated with the FCC and FAA have recently started experimenting with an implementation of the algorithm presented in this paper as a possible pricing mechanism for future auctions.

In this paper, we provide a new, more direct computational procedure for arriving at bidder-Pareto-optimal core outcomes in *any* sealed-bid combinatorial auction, using constraint generation and an explicit definition of the core region in payment space. Our technique may be viewed as an *approximate VCG mechanism*, and we provide concrete justification for the use of these mechanisms in general. At the crux of our approach, we show how to overcome one of the pitfalls of a linear programming (LP) approach to payment determination in the core, using constraint generation to handle the exponential number of constraints necessary to define the core. To accomplish this, we formulate the *core separation problem*, finding the most violated core constraint (most upset coalition) for any proposed payment vector. This core separation problem is also \mathcal{NP} -hard whenever winner determination is \mathcal{NP} -hard, and we show how to integrate the separation technique into a procedure that settles on a bidder-Pareto-optimal core point, with no need for limiting (substitutability) assumptions on the preferences of the bidders. Along the way, we solidify the concept of a *coalitional contribution*, measuring a winning bidder’s desire to join a coalition at a particular payment vector. We demonstrate the flexibility of our technique by showing a

few variations, and discuss the selection of a particular core outcome when several meet the “bidder-Pareto-optimal” criterion. In particular, we show that a mechanism that minimizes total payments consequently minimizes the total availability of gains from unilateral strategic manipulation, and is immune to a certain form of group collusion which would be profitable in any mechanism without this property. Finally, we compare the computational performance of our technique to others in the literature with a few detailed examples, and discuss directions for future research.

We begin the next section with the introduction of our notation, followed by a motivating example. In §3, we develop a general theory of bidder-Pareto-optimal core mechanisms to support our approach. In §4, we develop a specific algorithm for determining bidder-Pareto-optimal core payments, with comparisons to existing techniques in §5. Further, in §5 we reinforce our arguments on the selection of a bidder-Pareto-optimal payment vector by showing discrepancies among various solution techniques and by demonstrating a form of collusion that is nullified by our selection. We also provide a brief summary of our experience solving instances from the CATS data set. These comparisons demonstrate the effectiveness of the algorithm for computing bidder-Pareto-optimal payments presented in §4. Finally, in §6 we provide concluding remarks.

2. Problem Framework and Notation

Consider an environment where N distinct items are to be auctioned among M bidders. We will typically index each item in the set $I = \{1, 2, \dots, N\}$ of auction items by the letter i , and each bidder in the set $J = \{1, 2, \dots, M\}$ of all bidders by the letter j . For any set $S \subseteq I$, let $v_j(S)$ denote bidder j ’s value for the bundle of items S (the maximum amount he would be willing to pay for S), and let $b_j(S)$ denote the bid that bidder j has submitted to the auction for the bundle S . To maintain full generality, we assume an *XOR* language throughout (i.e., every bidder submits an exclusive bid for every possible bundle S), although the technique we propose in §4 easily generalizes to any bidding language that uses an IP formulation for winner determination. Further, we will use b_j and v_j to denote the full reports and valuations of bidder j over all bundles $S \subseteq I$.

To find an allocation of the auction items that maximizes bid value, we will make use of the following IP formulation for a general version of the winner determination (GWD) problem:

$$\begin{aligned} \max \quad & \sum_{j \in J} \sum_{S \subseteq I} b_j(S) \cdot x_j(S) && \text{(GWD)} \\ \text{subject to} \quad & \sum_{S \supseteq \{i\}} \sum_{j \in J} x_j(S) \leq 1 \quad \forall i \in I, && (1) \end{aligned}$$

$$\sum_{S \subseteq I} x_j(S) \leq 1 \quad \forall j \in J, \quad (2)$$

$$x_j(S) \in \{0, 1\} \quad \forall S \subseteq I, \forall j \in J. \quad (3)$$

The solution of this problem finds the efficient allocation for a given set of bids, represented by the coefficients $b_j(S)$ in the objective function. In this formulation, each binary variable $x_j(S)$ equals one if and only if bidder j is awarded bundle $S \subseteq I$. (Constraints (3) tell us that these variables must be binary.) Constraints (1) consequently ensure that each item is assigned to at most one bidder, while constraint set (2) prevents the auctioneer from accepting multiple bids from the same bidder. With “demand” constraints formulated as in (2), each bid is an exclusive offer made by a bidder for a particular bundle, and may not be combined or accepted in conjunction with any other offer from that bidder. The implied *bidding language* is therefore often referred to as an *XOR language*.²

Ausubel and Cramton (1999) make a strong and compelling case for the use of a mechanism (such as this one) that is efficient with respect to submitted bids, showing that it is revenue maximizing if the auction items can be costlessly resold. We also note that efficiency with respect to submitted bids contributes to our objective of perceived fairness in a public sector setting.

After establishing a winning allocation through the solution of an instance of GWD, we must also assign payments to bidders for the bundles they will receive. The variable π will be used to denote a payment vector, with each component π_j indicating the payment made by bidder j , and with superscripts used to distinguish among different payment vectors. For example, π_j^{VCG} will indicate bidder j 's VCG payment, while π^t will indicate the vector of payments for all bidders at iteration t in our payment adjustment procedure. In Theorem 3.2, notation of the form $\pi(b_j, b_{-j})$ will be used to denote the payment vector determined after a report of b_j by bidder j and a fixed set of bids b_{-j} by her competitors.

An *outcome*, Γ , refers to an allocation and set of payments for bidders in a combinatorial auction. Let the *coalition*, C_Γ , refer to the set of bidders receiving items under outcome Γ .

Perhaps the most obvious problem with VCG payments in terms of ex-post satisfaction is that the VCG

outcome may not be a “core outcome.” The core conditions for a one-sided (forward) auction can be stated by the following definitions, where a bidder weakly prefers outcome Γ_1 to Γ_2 if outcome Γ_1 gives him utility greater than or equal to the utility of outcome Γ_2 . We assume quasilinear net utility throughout (i.e., utility of a bundle S to bidder j is simply $v_j(S) - \pi_j$).

DEFINITION. An outcome Γ is *blocked* if there is an alternative outcome Γ_B which generates strictly more revenue for the seller and for which every bidder in C_{Γ_B} weakly prefers Γ_B to Γ . C_{Γ_B} may be referred to as a *blocking coalition*.

DEFINITION. An outcome Γ that is not blocked is called a *core outcome*.

DEFINITION. A core outcome Γ is *bidder-Pareto optimal* if there is no other core outcome weakly preferred by every bidder in C_Γ .

3. Bidder-Pareto-Optimal Core Mechanisms

To motivate the specific algorithm developed fully in §4, this section develops some general properties of sealed-bid mechanisms that produce efficient and bidder-Pareto-optimal core outcomes based on submitted bids.

We take as our primary motivation the well-known fact that a VCG outcome may not be a core outcome, and thus may be considered “socially unacceptable,” and unfit for an auction of public goods. Consider the following three-bidder, two-item example from Ausubel (2006): Let $b_1(AB) = b_2(A) = b_3(B) = 2$. In the efficient allocation, bidder 2 wins item A while bidder 3 wins item B . VCG payments for each winning bidder are computed to be zero. This is not a core outcome because bidder 1 would prefer to pay any amount up to two units to receive both items, an outcome which is clearly more desirable to the seller than the VCG outcome, in which no payments are collected. Even worse, if bidders 2 and 3 value the items at $v_2(A) = v_3(B) = 0.5$, they can still place bids of two units each, and under the VCG mechanism will not be held financially accountable for displacing the efficient winner, bidder 1. If instead bidders 2 and 3 each pay one unit given the bids of two units each, the coalition containing just bidder 1 no longer blocks, and the outcome is in the core. Further, the joint deviation from truthful bidding that benefitted bidders 2 and 3 when their valuations were each less than one is no longer profitable. Given this seemingly fatal flaw of the VCG mechanism, it is important in a public sector auction to eschew VCG payments in favor of a more reasonable payment determination mechanism that will be guaranteed to arrive at a core outcome.

Breaking away from the VCG mechanism is often met with skepticism from auction theorists: What

²When the bidders are constrained in the number of XOR bids they can submit (by a polynomial function in N , as may often be the case due to computing memory constraints), this XOR winner-determination problem is \mathcal{NP} -hard. For the problem in which bidders always submit bids on all bundles, there exist winner-determination algorithms which grow exponentially in the number of items N , but polynomially in the size of input (which is also exponential in N).

strategic behavior can we expect from bidders in a nonincentive-compatible auction? Ausubel and Milgrom (2002) characterize the Nash equilibria exactly for their own ascending proxy auction, describing a set of stable outcomes for which no player can benefit through unilateral deviation. Theorem 3.1 restates their result in a more general setting, justifying the auction design approach presented here.

To present our result, we employ Ausubel and Milgrom's (2002) notion of a semisincere bidding strategy, in which a bidder bids her true valuation on every bundle, minus some constant amount (or bids zero where this difference is negative). Ausubel and Milgrom (2002) show that such a strategy is a best response (maximizes utility) when facing a bidder-Pareto-optimal core mechanism, regardless of the bid profile of the competitors (e.g., even when the competitors do not play a Nash equilibrium strategy). For a fixed profile of opponent's bids, other strategies do exist that give the bidder utility equal to the best semisincere strategy when using either a bidder-Pareto-optimal or VCG mechanism, but these strategies can deliver less utility when the opponents do not bid as expected. Intuitively, the semisincere strategies provide equal attention to all bundles, bidding value minus a fixed profit margin on each, while other strategies focus greater attention on a particular bundle (or bundles), which will result in regret if unexpected opponents' bids demonstrate that attention was focused on the wrong bundle. We therefore restrict our attention to semisincere strategies, as in Ausubel and Milgrom (2002).

THEOREM 3.1. *Consider a mechanism that determines an efficient allocation and bidder-Pareto-optimal payments within the core, based on the submitted bids. Let A be the winning coalition with payments π^A and allocating bundle S_j^A to bidder j for any bidder-Pareto-optimal point in the core based on the true valuations of the bidders. An outcome in which every bidder bids $b_j(\cdot) = \max(v_j(\cdot) - v_j(S_j^A) + \pi_j^A, 0)$ is a Nash equilibrium under the given mechanism. Further, when using such a mechanism, any Nash equilibrium in semisincere strategies achieves a bidder-Pareto-optimal core point with respect to the true valuations of the bidders.*

PROOF. See the online appendix (provided in the e-companion).³ □

Theorem 3.1 tells us the behavior that can be expected in an environment of perfect information, demonstrating that for the expected equilibrium behavior in such an environment, the core property is obtained by using a mechanism that enforces

this property based on the submitted bids. Because the concept of a Nash equilibrium describes behavior under the possibility of profitable unilateral deviations (and not profitable group deviations), Theorem 3.1 provides only a broad set of possible outcomes (any of the bidder-Pareto-optimal points in the core) which may be expected in the case of perfect information. Thus, it is hard to say based on Theorem 3.1 alone which of the (sometimes many) possible Nash equilibria bidders will coordinate to achieve. We will show in §5, however, that under perfect knowledge and the possibility of hidden side payments to support group collusion, bidders can profitably coordinate to a *total-payment-minimizing* bidder-Pareto-optimal core point, narrowing the field of possible outcomes further (although not necessarily to a unique outcome).

Theorem 6 in Ausubel and Milgrom (2002) states that if the VCG outcome is in the core, then it is the only outcome that could be selected by such a bidder-Pareto-optimal core mechanism. Even though we have not fully specified all the details of the mechanism we intend to employ, we mention here that based on Theorem 3.1, together with Theorem 6 from Ausubel and Milgrom (2002), any mechanism that computes bidder-Pareto-optimal core payments based on the submitted bids dominates the VCG mechanism when core stability is a true concern. Whenever the VCG mechanism is well behaved (in the core), any bidder-Pareto-optimal mechanism will perform identically, and whenever the VCG outcome is poorly behaved (not in the core), a bidder-Pareto-optimal mechanism produces a better (i.e., core) outcome at equilibrium.

In real-life applications, however, information about competitors is not perfect, and may be costly to obtain or approximate. The expected cost of obtaining the information necessary to coordinate to an equilibrium may indeed be extremely high, when one considers that it may be illegal to obtain or share the pertinent information about one's competitors, and that violations of these legal restrictions result in exorbitant penalties. (The FCC, for example, periodically posts notices of such penalties on their website, www.fcc.gov.) What behavior should we expect in an environment with such a high cost of obtaining information about one's competitors?

As a first approximation of such an environment, we introduce the following two-player auction game. Suppose that players 1 and 2 can obtain utilities u_j (for $j = 1, 2$, respectively) by playing a particular Nash equilibrium strategy as described by Theorem 3.1. They do not, however, have enough information to compute the proper bids initially because without the knowledge of their competitor's valuations they cannot compute a bidder-Pareto-optimal core point.

³ An electronic companion to this paper is available as part of the online version that can be found at <http://manscijournal.informs.org/>.

Let us suppose that each may obtain the necessary information (including knowledge of her opponent’s decision to make a similar investment) at cost c_j . Implicitly, c_j includes the expected value of a penalty for violation of the auctioneer’s anti-collusion rules. Further, let r_j denote the reward obtained by bidder j if the other bidder tells the truth while j deviates from truth-telling according to her knowledge bought at cost c_j , while p_j denotes the penalty (extra payment) incurred by bidder j when telling the truth while the other bidder deviates.

Suppose next that the auctioneer in this situation chooses a mechanism that computes the same bidder-Pareto-optimal core outcome (when honest bids are reported) as the Nash equilibrium that the bidders coordinate to when both are able to obtain perfect information.⁴ The two bidders in this situation find themselves playing the following normal form game, in which each player has to decide whether to simply tell the truth (T) or pay the cost (F) to obtain the information necessary to find a more profitable outcome (let player 1 be the row player):

| | T | F |
|---|--------------------------------|--------------------------------|
| T | (u_1, u_2) | $(u_1 - p_1, u_2 + r_2 - c_2)$ |
| F | $(u_1 + r_1 - c_1, u_2 - p_2)$ | $(u_1 - c_1, u_2 - c_2)$ |

Like the prisoner’s dilemma, the two players will do best as a group if both play T, but for low values of c_j it will be a unique (pure-strategy) Nash equilibrium for both to play F, despite this being a worse outcome for both individually. However, if the values of c_j are made to be higher than the values of r_j and p_j (by the auctioneer enforcing large enough anti-collusion penalties), the unique Nash equilibrium becomes for both players to tell the truth. The bidders in this case are rewarded with the best possible outcome that they could coordinate to achieve, without having to pay the cost of obtaining the information necessary for such coordination. In an environment of more than two bidders, we obtain a similar result (concerning unilateral deviation) by letting this normal form game represent the problem facing a single bidder and allowing player 2 to represent the collection of all other bidders.

This suggests that the equilibrium behavior expected in a few extreme cases results in outcomes

⁴We assume that with perfect information comes perfect knowledge as to which of the (sometimes many) Nash equilibria should be coordinated to achieve. If this cannot be decided with perfect information, the situation becomes worse for the bidders, who may accidentally shade their bids to a noncore point and miss their optimal package/payoff. This possibility would only add benefit (safety) to the truthful strategy outcome, and does not drastically alter our discussion.

with the desired core property. When information is costlessly available, bidders will coordinate to a Nash equilibrium via Theorem 3.1. When information is extremely expensive, we expect the bidders to tell the truth and allow the mechanism to prescribe a satisfactory outcome as in the two-player game above. In between (when the cost of information is moderate), bidders may be expected to needlessly invest in perfect information, being unable to trust their competitors not to fink, as in the traditional prisoner’s dilemma. In all cases, however, the expected behavior results in a core outcome.

We note that our model treats the decision of investment in information as a 0–1 decision and that a future analysis of partial investment for partial information remains an interesting open problem. We conjecture that such analysis will verify the following intuition: bidders will “shade” their bids as much as possible with “cheap” information (e.g., the information that is privately available to them without illicit pre-auction information sharing), resulting in the selection of a bidder-Pareto-optimal core point by the convexity of the core. With equal shading ability, this results in the same outcome as under honest revelation, but otherwise, those who shade better will be rewarded with a more preferred bidder-Pareto-optimal core point.

The analysis thus far suggests the following design objectives for an auction mechanism:

- Employ a mechanism that computes a bidder-Pareto-optimal core outcome based on submitted bids, a choice that dominates the use of a VCG mechanism when core stability is desired.
- Among all bidder-Pareto-optimal core outcomes, select one that the bidders will naturally coordinate to in the presence of perfect information. This will diminish the need for investment in information about one’s competitors, which may be illegal. (More on the selection of an outcome will be discussed as we proceed.) When information is costly, this will make truth-telling a preferred outcome for every bidder relative to the outcome in which all fully invest in information.
- Enforce harsh penalties for collusion (the illegal sharing of information). We do not address the issue of just how large these penalties should be, although we have made a case for all bidders benefitting (and the seller being indifferent) from a scenario in which such penalties are arbitrarily large. In practice, a political process typically determines the magnitude of these penalties for public sector auctions.
- Make the reward for unilateral deviation from truth-telling (in the case where all others tell the truth) as small as possible. This helps encourage the situation of extremely costly information (and not moderately costly information), in which costs of deviation are high relative to the benefits of unilateral deviation.

To address this last point, Theorem 3.2 provides a bound on a bidder’s ability to benefit by unilaterally deviating from the truthful reporting strategy given that others each play a truthful strategy. In particular, the difference between a bidder’s payment using the VCG mechanism and an alternative mechanism (assuming truthful reports) is the maximum that a bidder can gain by deviating in the alternative mechanism. Using the corollary that follows, these bounds provide us with the ability to minimize the total gains from deviation (over all bidders). These results provide support for an approximate VCG approach in the most intuitive fashion: the closer you get to the VCG payments, the less incentive there is to deviate.

Suppose that each bidder j reports a vector b_j , and that instead of assigning VCG payments $\pi_j^{\text{VCG}}(b_j, b_{-j})$, our auction mechanism uses a payment rule $\pi(b_j, b_{-j})$ that is guaranteed to assign payments greater than or equal to the VCG payments. (This assumption is justified because core payments are always greater than or equal to the VCG payments.) Also, assume that the mechanism chooses the efficient allocation (based on the reports), so that for the same reports the mechanism in question differs from the VCG mechanism in payments only (i.e., allocations are the same under both mechanisms for the same reports). We then have the following result:

THEOREM 3.2. *For any efficient mechanism that produces payments π_j greater than or equal to the VCG payments, the amount that bidder j can benefit by unilaterally deviating from the honest report strategy is less than or equal to $\pi_j(v_j, b_{-j}) - \pi_j^{\text{VCG}}(v_j, b_{-j})$.*

PROOF. Suppose not: There is some report \hat{b}_j such that

$$(v_j(\hat{S}_j) - \pi_j(\hat{b}_j, b_{-j})) - (v_j(S_j) - \pi_j(v_j, b_{-j})) > \pi_j(v_j, b_{-j}) - \pi_j^{\text{VCG}}(v_j, b_{-j}),$$

where \hat{S}_j is the bundle awarded to bidder j given the reports (\hat{b}_j, b_{-j}) , and S_j is the efficiently awarded bundle given the report (v_j, b_{-j}) . After rearranging and canceling, we have

$$v_j(\hat{S}_j) - \pi_j(\hat{b}_j, b_{-j}) > v_j(S_j) - \pi_j^{\text{VCG}}(v_j, b_{-j}).$$

By assumption on the mechanism determining π (and a standard assumption of quasilinear utility), we have

$$v_j(\hat{S}_j) - \pi_j^{\text{VCG}}(\hat{b}_j, b_{-j}) \geq v_j(\hat{S}_j) - \pi_j(\hat{b}_j, b_{-j}),$$

and thus from the above two inequalities, we find

$$v_j(\hat{S}_j) - \pi_j^{\text{VCG}}(\hat{b}_j, b_{-j}) > v_j(S_j) - \pi_j^{\text{VCG}}(v_j, b_{-j}).$$

But this contradicts the well-known incentive-compatibility property of the VCG mechanism, thus the supposition must be false. \square

This theorem and its elementary proof lend theoretical support to several payment-setting techniques for combinatorial auctions (or two-sided exchanges) which eschew the VCG mechanism, but approximate the VCG outcome as closely as possible within a feasible set of payments. These bounds are first stated in Parkes et al. (2001) in the context of a two-sided exchange, who also demonstrate that they are tight (that is, by bidding exactly the VCG payment, a bidder can always achieve this benefit, given that all other bidders remain honest). We provide a rigorous proof of these bounds in the context of a one-sided auction and use them to justify the approach of approximating VCG payments for one-sided auctions.

For example, suppose that we want to determine a payment rule that assigns payments to winning bidders in an efficient allocation, and additionally, that these payments must satisfy some prespecified properties (for example, the payments are not blocked by any coalition of bidders). If R is the region in payment space containing all payment vectors satisfying the prespecified properties (for example, the core), assuming that R is bounded below by the VCG payments, we have the following result:

COROLLARY 3.3. *The payment rule that minimizes total potential gains from deviation within the feasible region R minimizes total payments over R .*

PROOF. When minimizing the sum of the (tight) bounds provided by Theorem 3.2, the VCG payment terms are constant relative to the payments, with the result following. \square

In Parkes et al. (2001), VCG payments were approximated subject to budget balance in a two-sided combinatorial exchange. The authors investigated several payment rules for VCG approximation, eventually favoring a “threshold rule” in which payments are selected that minimize the maximum difference from VCG payments over all players. In the case of a budget-balanced exchange, however, total payments are constant (equal to zero), and thus Corollary 3.3 has no bite. In the case of a one-sided auction in which VCG payments are approximated within the core, one approach may be to simply follow the lead of Parkes et al. (2001) and attempt to find payments that minimize the maximum difference from VCG payments within the core. Corollary 3.3, however, shows that we must first minimize total payments, or else the global ability to gain from deviation may not be minimized. In §5, we provide an example demonstrating that a set of payments that minimizes the maximum difference from the VCG payments may not minimize total payments, and our algorithm developed in §4 shows how to incorporate the minimization of the maximum deviation from VCG as a secondary objective.

Having established that our goal is to minimize total payments within the core region in payment space, we now explore the problem of properly defining the core using linear and integer programming. Later we will see that the minimization of total payments is also beneficial with regards to our second bulleted design objective (above) when one considers the possibility of profitable group deviation.

4. Core Constraint Generation

We now provide a novel contribution, an algorithm for determining core payments using constraint generation. Our technique requires the repeated use of a slightly modified winner-determination problem, for which GWD will serve as a starting point. Although the XOR language may seem cumbersome to those who appreciate the benefits of a more sophisticated bidding language, we present our results in the XOR language because it is fully expressive, allowing for the expression of any demand function through complete enumeration, and therefore most general. The method we present, however, may be applied in the context of any bid language, by simply performing the technique introduced in the following section on the corresponding winner-determination problem for that particular bidding language.

As in Hoffman et al. (2006), our general approach is to first solve several winner-determination problems for a sealed-bid combinatorial auction, settling on a particular set of winning bidders and VCG payments. Given this fixed outcome Γ_{VCG} , we denote the associated coalition of winners as $W = C_{\Gamma_{VCG}}$ and the associated payments π_j^{VCG} for each bidder $j \in W$. Starting from this VCG outcome, we wish to arrive at a bidder-Pareto-optimal core outcome with the same efficient allocation; for the same bids only payments (if anything) will differ between the VCG outcome and the final core outcome proposed here.

From a constrained optimization point of view, it is difficult to categorize the bidder-Pareto-optimal core payments for two reasons. First, an exponential number of possible coalitions must be considered to define the core region in payment space; each coalition (subset of the set J of all bidders) is willing to make some offer to the seller, and the total payments of the winning bidders must exceed each of these offers. Second, it is difficult to gauge how much a winning bidder will contribute to a “coalitional value function” (the offer made to the seller) without knowledge of his final payment.

Given the core definition provided in §2, it seems convenient (at first) to define a core constraint for any blocking outcome Γ as

$$\sum_{j \in W} \pi_j \geq z_C,$$

where π_j is a payment made by each bidder $j \in W$, and z_C is the coalitional value of C . We define this coalitional value z_C as the maximum total payments that a coalition C would be willing to offer the seller in any outcome Γ with $C_\Gamma = C$. Noting that an exponential number of such coalitions exist, each defining a constraint on the acceptable outcomes in payment space, how do we separate which constraints must be applied from those that can be ignored without consequence? Further, how do we compute the value of z_C for a particular coalition C ?

To appreciate the subtlety of the second question, observe that any bidder $j \in W$ would not want to join a coalition and receive less surplus than he would at his current payment π_j^t , at any iteration of a price-adjustment algorithm or point in time t . Trying to solve the winner determination problem restricted to the bidders in the coalition C therefore overstates the amount that any winning bidder in C would contribute to the coalitional value z_C because he will not be compensated his opportunity cost (the amount of surplus he stands to gain at the current vector of payments). The result would be the application of a too restrictive constraint, therefore charging the winning bidders too much. We must instead consider that each winning bidder has an opportunity cost limiting his contributions to a blocking coalition. We therefore amend our notation to capture the fact that a coalitional value function depends on the current set of payments π^t , and thus change to $z_C(\pi^t)$.

We may now describe our general iterative approach as follows: at a current vector in payment space, π^t at each discrete iteration t , find the coalition C^t with the highest coalitional value relative to current payments, $z_C(\pi^t)$. If this coalition blocks the currently proposed allocation (the efficient solution with payments π^t), apply the corresponding core constraint and find a new bidder-Pareto-optimal payment π^{t+1} . If this coalition does not block, then terminate at the currently proposed outcome. Before demonstrating this technique in full detail, we observe a few basic properties of each winning bidder’s “coalitional contribution” relative to current payments, leading us to a more useful LP characterization.

Let S_j be the bundle awarded to winning bidder j in the efficient solution. Relative to the current payment vector, bidder j would not voluntarily join a coalition that offers him less surplus than $b_j(S_j) - \pi_j^t$, his opportunity cost. If a coalition provides bidder j with bundle \bar{S}_j (possibly the same as S_j , but in general not), then bidder j would contribute at most $b_j(\bar{S}_j) - (b_j(S_j) - \pi_j^t)$ into the coalition value function; if he contributed more, then he would perceive less benefit (surplus) from the hypothetical outcome of the coalition than the one proposed by the auctioneer at iteration t . If bidder j is not in the set W , there is no

opportunity cost to recover, and he would offer his entire value for a bundle in an effort to block the set of winning payments. In this case, $b_j(S_j = \emptyset) = \pi_j^t = 0$ and his coalitional contribution would be $b_j(\bar{S}_j)$. In general, we have the following definition:

DEFINITION. Facing payment of π_j^t for bundle S_j in the efficient allocation, bidder j would be willing to make a *coalitional contribution* of $q_j(\bar{S}_j, \pi_j^t) = b_j(\bar{S}_j) - b_j(S_j) + \pi_j^t$ to receive \bar{S}_j as part of the coalition C .

LEMMA 4.1. We note the following basic properties of the coalitional contribution $q_j(\bar{S}_j, \pi_j^t)$:

(1) If bidder j were to pay an amount π_j^c that is greater than $q_j(\bar{S}_j, \pi_j^t)$ to join coalition C and win \bar{S}_j , then bidder j would experience less surplus from the coalition than from the auctioneer's proposed outcome S_j with payment π_j^t .

(2) $q_j(\bar{S}_j, \pi_j^t)$ increases linearly in $\pi_j^t \forall j \in W$.

(3) $\forall j \in J, \bar{S}_j = S_j$ implies that $q_j(\bar{S}_j, \pi_j^t) = \pi_j^t$.

(4) $q_j(\bar{S}_j, \pi_j^t) = b_j(\bar{S}_j)$ and $j \in W \cap C$ imply that $\pi_j^t = b_j(S_j)$.

(5) $q_j(\bar{S}_j, \pi_j^t) = b_j(\bar{S}_j) \forall j \notin W$.

PROOF. Property (1) verifies that our definition of the quantity $q_j(\bar{S}_j, \pi_j^t)$ truly reflects what we mean by coalitional contribution, and follows from the definition: $\pi_j^c > q_j(\bar{S}_j, \pi_j^t) = b_j(\bar{S}_j) - b_j(S_j) + \pi_j^t$ implies that surplus from the coalitional outcome is $b_j(\bar{S}_j) - \pi_j^c < b_j(S_j) - \pi_j^t$, where the right-hand side of the last inequality is exactly the surplus from the auctioneer's outcome. Properties (2), (3), and (4) follow directly from the definition, while Property (5) follows from the standard assumptions that $b_j(\emptyset) = 0 \forall j$, and that $\pi_j^t = 0 \forall j \notin W$. \square

Now that we have established the importance of opportunity cost and shown that bidder j 's coalitional contribution will normally be less than his value for the bundle given to him by the coalition, we formulate the *core-constraint separation problem* at payment vector π^t . At any point in payment space, π^t , the integer program SEP^t finds the most violated core constraint, or tells us that no blocking coalition can be found:

$$z(\pi^t) = \max \sum_{j \in J} \sum_{S \subseteq I} b_j(S) \cdot x_j(S) - \sum_{j \in W} (b_j(S_j) - \pi_j^t) \cdot \gamma_j \quad (SEP^t)$$

$$\begin{aligned} \text{subject to } & \sum_{S \supseteq \{i\}} \sum_{j \in J} x_j(S) \leq 1 \quad \forall i \in I, \\ & \sum_{S \subseteq I} x_j(S) \leq 1 \quad \forall j \in J \setminus W, \\ & \sum_{S \subseteq I} x_j(S) \leq \gamma_j \quad \forall j \in W, \\ & x_j(S) \in \{0, 1\} \quad \forall S \subseteq I, \forall j \in J, \\ & \gamma_j \in \{0, 1\} \quad \forall j. \end{aligned}$$

The added terms in the objective of this IP formulation (starting from GWD in §2) ensure that any winning bidder will be compensated his opportunity cost if selected as part of the optimal solution to SEP^t . It is easy to verify that any bidder $j \in J$ contributes exactly $q_j(\bar{S}_j, \pi_j^t)$, his coalitional contribution, and that this formulation is equivalent to an instance of the winner-determination problem for the auction in which each bid from a winning bidder is reduced by his current opportunity cost.

If the objective $z(\pi^t) > \sum_{j \in W} \pi_j^t$, then the bidders for which $x_j(S) = 1$ form a coalition C^t which blocks the efficient allocation with the current set of payments, i.e., C^t is a set of bidders such that $z(\pi^t) = z_{C^t}(\pi^t)$. If $z(\pi^t) = \sum_{j \in W} \pi_j^t$, then we have achieved an unblocked core outcome; no coalition of bidders would be able to offer an outcome that both they and the auctioneer would prefer. Also note that $z(\pi^t) < \sum_{j \in W} \pi_j^t$ is not a possibility because the feasible allocation of items to the winning set W achieves an objective value of $\sum_{j \in W} \pi_j^t$. As a result, the algorithm developed in this paper is ascending in terms of total payments.

Having found a coalition C^t that blocks the proposed outcome with payment vector π^t , we know that the constraint $\sum_{j \in W} \pi_j \geq z_{C^t}(\pi^t)$ is violated. We could proceed to define the core using constraints of this form, however, Property (2) from Lemma 4.1 indicates that any price increases on any bidder $j \in W \cap C$ completely cancel with the corresponding increase in coalitional contribution $q_j(\bar{S}_j, \pi_j^t)$ on the right-hand side of the constraint for C^t . Consequently, our method is to find bidder-Pareto-optimal core payments by generating core constraints of the form $\sum_{j \in W \setminus C^t} \pi_j \geq z(\pi^t) - \sum_{j \in W \cap C^t} \pi_j^t$. The right-hand side of these constraints are computed once and remain constant for the remainder of the algorithm, unlike the right-hand side of the constraints in $\sum_{j \in W} \pi_j \geq z_{C^t}(\pi^t)$.

After finding each of these constraints using SEP^t , we may solve the following linear program BPO^t to find a set of bidder-Pareto-optimal core payments, in the core relative to all coalitions found through iteration t . Our choice of objective function is motivated by Corollary 3.3 (and is reinforced later by an anti-collusion property in §5).

$$\theta^t = \min \sum_{j \in W} \pi_j \quad (BPO)^t$$

$$\begin{aligned} \text{subject to } & \sum_{j \in W \setminus C^\tau} \pi_j \geq z(\pi^\tau) - \sum_{j \in W \cap C^\tau} \pi_j^\tau \\ & \forall \tau \leq t, \quad (\text{CORE}) \end{aligned}$$

$$\pi_j^{\text{VCG}} \leq \pi_j \leq b_j(S_j) \quad \forall j \in W.$$

We then use the value of each π_j in the solution for the next iteration (i.e., set $\pi_j^{t+1} = \pi_j$). We now show that the resulting ascending algorithm converges to the desired outcome.

THEOREM 4.2. *If $z(\pi^t) = \theta^{t-1}$, then the solution to BPO^{t-1} yields bidder-Pareto-optimal core payments.*

PROOF. First, we note that the payments π^t are in the core because if not a blocking coalition provides a solution to SEP^t greater than θ^{t-1} . By the minimality of $\sum_{j \in W} \pi_j$ with respect to the constraints (CORE), if any bidder experiences a payment reduction without an increase in the payment of some other bidder, then some constraint from (CORE) must be violated. Thus, we find that the payments must be Pareto efficient with respect to core constraints. The convergence of this algorithm is guaranteed because only a finite number of constraints may be generated, and because the region (CORE) always contains at least the trivial pay-as-bid solution. \square

In some cases, there may be multiple optimal solutions to BPO^t for which we suggest the following refinement, a linear program $EBPO^t$ which finds equitable bidder-Pareto-optimal core payments:

$$\begin{aligned} \theta^t(\varepsilon) = \min \quad & \sum_{j \in W} \pi_j + \varepsilon m && \text{(EBPO}^t\text{)} \\ \text{subject to} \quad & \sum_{j \in W \setminus C^\tau} \pi_j \geq z(\pi^\tau) - \sum_{j \in W \cap C^\tau} \pi_j^\tau \quad \forall \tau \leq t, \\ & \pi_j - m \leq \pi_j^{\text{VCG}} \quad \forall j \in W, \\ & \pi_j \leq b_j(S_j) \quad \forall j \in W, \\ & \pi_j \geq \pi_j^{\text{VCG}} \quad \forall j \in W, \end{aligned}$$

with decision variables appearing only on the left-hand side of each constraint. By taking a small enough value of the scalar ε , we have that $\sum_{j \in W} \pi_j$ computed from the solution to this problem is equal to θ^t from the corresponding instance of BPO^t . We then see that the effect of the new terms is to find a set of payments that minimizes the maximum difference from the VCG payments over all bidders, among all outcomes that minimize the total payments of winning bidders. The entire process for determining equitable bidder-Pareto-optimal core payments is summarized in Table 1.

Because a general instance of the winner-determination problem can be polynomially transformed to

Table 1 The Core Constraint Generation Algorithm

| |
|--|
| Initialize: Iteratively solve the winner-determination problem, finding winners W and VCG payments; Set $\pi^1 = \pi^{\text{VCG}}$, and $\theta^0(\varepsilon) = \sum_{j \in W} \pi_j^{\text{VCG}}$. |
| Iteration t : Solve SEP^t |
| If $z(\pi^t) > \theta^{t-1}(\varepsilon)$ |
| Add constraint $\sum_{j \in W \setminus C^t} \pi_j \geq z(\pi^t) - \sum_{j \in W \cap C^t} \pi_j^t$ to $EBPO^t$ and solve, updating $\theta^t(\varepsilon)$. |
| Set $\pi_j^{t+1} = \pi_j$, the values found from $EBPO^t$. |
| Iterate: $t = t + 1$. |
| Else |
| Terminate. |

the problem SEP^t (by adding a fictional bidder who wins everything and solving the separation problem when this bidder pays zero), the separation problem is \mathcal{NP} -hard whenever winner determination is \mathcal{NP} -hard. Consequently, by a well-known result establishing the equivalence of separation and optimization (see Grötschel et al. 1981), the problem of optimizing a linear objective over the CORE region is also \mathcal{NP} -hard (whenever winner determination is \mathcal{NP} -hard).

Thus, a variety of related payment rules are also fundamentally difficult to compute. Among these are a minimax deviation from the VCG payments rule (or threshold rule), and a maximin deviation from the bundle values (or equal-pay rule). Versions of these rules are discussed by Parkes et al. (2001) in the context of two-sided exchanges with budget balance. The modified objective function in $EBPO^t$ allows one to mimic the equity properties of the Parkes et al. (2001) threshold rule as a secondary objective, while minimizing the total ability of bidders to gain from unilateral deviation according to Corollary 3.3.


In the next section, we illustrate this core constraint generation algorithm with an example, and compare the results to those of the threshold rule, and to the examples presented in Hoffman et al. (2006) and Wurman et al. (2004) regarding proxy auction outcomes. In §5.1, we establish the primacy of total-payment minimization over the space of all bidder-Pareto-optimal solutions, and consider a related Nash bargaining game in §5.2.

5. Examples, Comparisons to Other Techniques, and Some Computational Experiments

To illustrate our technique for determining bidder-Pareto-optimal core outcomes using core constraint generation, consider the three-item, eight-bidder example of Figure 1. Here we depict the three items as pieces of a pie to indicate visually which bids are on which items. Inspection shows that bidders 1, 2, and 3 constitute the set of winners in the efficient allocation, as indicated. In accordance with our earlier notation, we will say that $W = \{1, 2, 3\}$.

With the set of winners determined, we next compute $\pi_1^{\text{VCG}} = \pi_2^{\text{VCG}} = \pi_3^{\text{VCG}} = 10$, and so set $\pi_1^1 = \pi_2^1 = \pi_3^1 = 10$. Solving SEP^1 , the separation problem at iteration 1, we find the most violated blocking coalition $C^1 = \{3, 4\}$ with $z(\pi^1) = 38 > 30 = \sum_{j \in W} \pi_j^1$. We then add the constraint $\pi_1 + \pi_2 \geq 28$ to the formulation $EBPO^1$ and solve to find the new set of payments $\pi_1^2 = \pi_2^2 = 14$, $\pi_3^2 = 10$. Next, we solve SEP^2 and find that the coalition $C^2 = \{2, 5\}$ blocks the current set of payments with $z(\pi^2) = 40 > 38 = \sum_{j \in W} \pi_j^2$. We form $EBPO^2$ by adding the constraint $\pi_1 + \pi_3 \geq 26$ and

Figure 1 An Auction for Which Threshold Payments Do Not Minimize Total Payments



solve, yielding $\pi_1^3 = 16$, $\pi_2^3 = 12$, and $\pi_3^3 = 10$. Finally, solving SEP³, we find that no blocking coalitions exist; the process has terminated at a bidder-Pareto-optimal core outcome.

5.1. The Threshold Rule and a Collusion Problem

The algorithm we have just demonstrated selects a specific bidder-Pareto-optimal core outcome. We now compare our method to the “minimax” rules for finding bidder-Pareto-optimal core outcomes that are analogous to the rules developed for finding balanced-budget solutions in combinatorial exchanges (Parkes et al. 2001). For example, reformulated in our own notation, the *threshold rule* selects payments which solve the following optimization problem:

$$\begin{aligned} & \min \max_{j \in W} (\pi_j - \pi_j^{\text{VCG}}) \\ & \text{subject to } \sum_{j \in W \setminus C} \pi_j \geq n_C \quad \forall C \subseteq J, \end{aligned} \quad \text{(THRESH)}$$

where n_C denotes the portion of the coalitional value for C that is attributable to nonwinners. This rule for payment selection can be computed as an LP:

$$\begin{aligned} & \min m \\ & \text{subject to } \sum_{j \in W \setminus C} \pi_j \geq n_C \quad \forall C \subseteq J, \\ & \pi_j - \pi_j^{\text{VCG}} \leq m \quad \forall j \in W. \end{aligned} \quad \text{(THRESH-LP)}$$

Solving this LP for the above example, we find the solution $\pi_1 = \pi_2 = \pi_3 = 14$, which is unfortunately Pareto inefficient. There is, however, a bidder-Pareto-optimal point among the optimal solutions to THRESH-LP, and we may next consider the strength of the threshold rule as a selection criteria among the multitude of points on the Pareto frontier of the core. As in our own formulation for selecting equitable payments, this may be accomplished with the insertion of a secondary objective weighted by a tiny

value ε :

$$\begin{aligned} & \min m + \varepsilon \cdot \sum_{j \in W} \pi_j \\ & \text{subject to } \sum_{j \in W \setminus C} \pi_j \geq n_C \quad \forall C \subseteq J, \\ & \pi_j - \pi_j^{\text{VCG}} \leq m \quad \forall j \in W. \end{aligned} \quad \text{(Pareto-THRESH)}$$

Solving this new LP for the example of this section, we arrive at the solution $\pi_1 = \pi_2 = 14$ and $\pi_3 = 12$, which is indeed Pareto efficient and minimizes the maximum deviation from the VCG payments. We note, however, that in accordance with the bounds presented in Theorem 3.2, this solution provides 10 units of opportunity to gain from deviation, while the EBPO^t solution provides only eight units of opportunity to gain from deviation, a global decrease in the opportunities from deviation, although individual bidders may have greater opportunities.

The selection of a bidder-Pareto-optimal outcome is by its very nature a matter of taste. Bidder-Pareto-optimal points are by definition unable to be compared with a strict dominance relationship; movement from one Pareto optimal point to another results in an increase in utility to one player if and only if another player experiences a decrease in utility. Although it seems reasonable to guarantee that total available incentive to deviate be minimized, one could argue that minimizing the maximum incentive to deviate is more important (the previous example demonstrates that there is a distinction). We note, however, the following important observation:

PROPOSITION 5.1. *A core outcome that minimizes the total payments by bidders will be strictly preferred by the bidders to any other bidder-Pareto-optimal core outcome if side payments are possible.*

According to Theorem 3.1, because collectively bidding any bidder-Pareto-optimal point in the core constitutes a Nash equilibrium, one must look outside the concept of unilateral deviation to compare bidder-Pareto-optimal outcomes. To select a payment mechanism from all bidder-Pareto-optimal choices, we must therefore consider multilateral collusion supported by side payments.

Compare, for example, the outcome arrived at by the solution of EBPO^t, $\pi_1 = 16$, $\pi_2 = 12$, $\pi_3 = 10$, and the outcome arrived at by the solution of Pareto-THRESH, $\pi_1 = \pi_2 = 14$, $\pi_3 = 12$. In the former case, the sum of the payments equals 38, while in the latter case, the payments sum to 40. If asked which outcome is preferred, in the presence of side payments the winners will select the former. This is because the bidders that stand to benefit from the change (bidders 2 and 3, in this case) can pay off the bidders who would experience a payment increase (just bidder 1) to make

them indifferent to the change. In this case, bidders 2 and 3 could compensate bidder 1 a payment of one unit each in an effort to coax him into a move from the latter to the former. Bidder 1 would then be indifferent, while bidders 2 and 3 will each still gain one unit of surplus from the change. We note that the same comparison holds for the “equal-pay rule” (in which the maximum difference between the bundle values and payments is minimized) because by the construction of the example in Figure 1 all winning bidders have the same valuations and VCG payments.

PROPOSITION 5.2. *For any auction mechanism terminating at a bidder-Pareto-optimal core outcome, it is profitable for the winning bidders to collude to bid (and pay) exactly the final payments prescribed under a total-payment-minimization mechanism (e.g., EBPO^t).*

This result is achieved by using covert side payments to support any difference from the mechanism selected by the auctioneer, as shown in the preceding example. Further, as in the preceding example, this collusion will be advantageous because total payments are reduced, leaving money available to support the covert side payments.

Interestingly, if the public were made aware of the side payments, and each bidder’s apparent final payment for her awarded bundle is adjusted accordingly, the resulting solution is no longer in the core. For example, if bidders facing the threshold rule make side payments to support a move to the EBPO^t solution for the Figure 1 auction as mentioned earlier, the adjusted final payments (adding or subtracting side payments) are $\pi_1 = 14$, $\pi_2 = 13$, $\pi_3 = 11$. But this outcome is blocked; bidder 4 would be willing to pay up to 28 monetary units for the items won by bidders 1 and 2 for just 27 monetary units.

Because this possibility of colluding to a noncore point exists whenever the auction mechanism does not minimize total payments within the core, we argue that the auction mechanism should necessarily minimize total payments within the core. By selecting an outcome that would be emulated under the collusive strategy, we eliminate the incentive to engage in such collusion with illicit payments.

5.2. Nash Bargaining

With arguments in the previous subsection proposing the selection of an outcome that winners may try to achieve through collusion under a different mechanism, we must naturally consider the connections to the theory of bargaining. In a seminal paper, Nash (1950) introduces a solution concept for a set of players bargaining over a convex set of feasible outcomes in payoff space. Under a few mild assumptions, it can be shown that the players should immediately settle on the point that maximizes the product of their util-

ities within this convex set. Because the convexity of the core is a well-known property, one may wonder if there is an interpretation of our auction mechanism as a Nash bargaining game, and whether the outcomes coincide. After all, we are proposing to select the outcome that bidders can profitably collude to under some other mechanism. Is this the same solution that Nash bargaining says these bidders will naturally agree to? Consider the following game:

Bidders are told that they are participating in a VCG auction and therefore respond with their true valuation for every bundle. The auctioneer releases the auction (VCG) outcome including the efficient winners and their payments, causing great discontent among losing bidders because the solution is not in the core. As a political solution to the public outcry, the auctioneer asks the winning bidders to decide among themselves on a set of payments that no group of dissatisfied bidders could complain about (i.e., a core outcome), or else the auction results will be canceled with no trade. This is a Nash bargaining game with the core as a feasible region. The solution is the point in the core that maximizes the product of the winning bidder utilities.

Is the solution of this bargaining game the same as the solution found using core constraint generation? For some examples yes, but in general the answer is no. Again, take the example in Figure 1. The unique solution to the Nash bargaining problem for this auction is $\pi_1 = 16 - \lambda$, $\pi_2 = 12 + \lambda$, $\pi_3 = 10 + \lambda$, where

$$\lambda = \frac{14 - 2\sqrt{43}}{3}.$$

Finding this solution requires the solution of a nonlinear programming problem, which in general may be quite difficult. It is easy to verify that the product of all winning bidders’ utilities is higher under this solution than the one found using core constraint generation. To verify that this is indeed the optimal solution to the Nash bargaining problem, one need only verify the pseudoconcavity of the objective function (the product of three quasilinear nonnegative utility functions) over the core, and that this point satisfies the Karush-Kuhn-Tucker optimality conditions (see, for example, Bazaraa et al. 1979). The derivation of this solution is a bit tedious, but would make a healthy exercise for a student of nonlinear programming.

Is this a better core solution for this auction than the one found using core constraint generation? We first note that it suffers from the same problem as the threshold solution (noted in §5.1) that total payments are not minimized (payments sum to $38 + \lambda$, rather than the minimal 38). Thus, if an auction mechanism selects the payment vector within the core that maximizes the product of apparent utility, there

will be a greater total incentive to deviate via Corollary 3.3. Further, taking into consideration the comments on side payments from §5.1, if the winning bidders were allowed to select a “socially acceptable” outcome within the core as depicted above, they would prefer to choose a payment minimizing outcome, with bidders 2 and 3 each covertly paying bidder 1 the amount $\lambda/2$. The bidders would announce the “socially acceptable” outcome $\pi_1 = 16$, $\pi_2 = 12$, $\pi_3 = 10$, but with side payments would be achieving the “socially unacceptable” outcome $\pi_1 = 16 - \lambda$, $\pi_2 = 12 + \lambda/2$, $\pi_3 = 10 + \lambda/2$. The Nash bargaining model does not take into consideration the expansion of the feasible region by use of side payments, nor does it consider the incentive properties of the auction giving rise to this feasible region, although it is interesting to note the connection and the disparity between the two solution concepts.

5.3. Accelerated Proxy Methods

Having established with the example of Figure 1 that the outcome of the core constraint generation procedure may deviate from the threshold outcome as formulated in Hoffman et al. (2006) and Parkes (2002), and from a related Nash bargaining game solution, we now demonstrate that this technique may provide computational advantages over existing iterative proxy methods. Consider the following example from Hoffman et al. (2006):

| Bidder | 1 | 2 | 3 | 4 |
|---------|----|-----|----|----|
| Package | AB | BC* | AC | A* |
| Value | 20 | 26 | 24 | 16 |

where * in their notation denotes the winners in the efficient allocation. They present a comparison of several methods for obtaining core outcomes, and here we present their results for this problem in terms of number of rounds, along with the VCG and threshold payments, with an added row displaying the new results using Core Constraint Generation. Table 2 shows that where each of their iterative techniques⁵ may require several rounds to solve this problem, core constraint generation terminates after a single round of price adjustments. Starting at the VCG payments, $\pi_2^{VCG} = 8$, $\pi_4^{VCG} = 0$, we find the most violated blocking coalition consisting of just bidder 3. We then equitably divide the burden of overcoming this coalition

⁵ “Pure proxy” refers to the algorithm described by Ausubel and Milgrom (2002) (increment = 0.01), while “safe start” runs this same algorithm starting from the VCG payments. “Increment scaling” uses a similar algorithm with a changing increment value, starting either from zero payments or from VCG (i.e., w/ “safe start”), proposed by Hoffman et al. (2006), who give full details on these methods.

Table 2 Comparison of Core Constraint Generation to Proxy Methods

| Method | Rounds | Revenue | Payment by 2 | Payment by 4 |
|---------------------------------|--------|---------|--------------|--------------|
| Pure proxy | 3,100 | 24.02 | 12.01 | 12.01 |
| Safe start | 800 | 24.02 | 16.01 | 8.01 |
| Increment scaling | 20 | 24.02 | 17.01 | 7.01 |
| Increment scaling w/ safe start | 15 | 24.02 | 16.01 | 8.01 |
| VCG payments | — | 8.00 | 8.00 | 0.00 |
| Threshold payments | — | 24.00 | 16.00 | 8.00 |
| Core constraint generation | 1 | 24.00 | 16.00 | 8.00 |

(using EBPO^t) and find that no other blocking coalition exists. Note that this procedure obviates the need to consider a constraint for the coalition {1} which is made redundant by the constraint of coalition {3}.

Comparisons to all other examples worked out by Hoffman et al. (2006) verify this apparent dominance of the core constraint generation technique. For every problem instance presented fully there, core constraint generation terminates after a single price adjustment. We also note that in no case presented in Hoffman et al. (2006) do our computed payments differ from the threshold payments, and conclude that the phenomenon present in the example of Figure 1 is missing from the examples generated by Hoffman et al. (2006).

5.4. The Inflection Point Method

A different technique for direct computation of proxy outcomes (i.e., bidder-Pareto-optimal core outcomes) is provided by Wurman et al. (2004), who observe the existence of change points or inflection points in the price trajectories followed in the typical proxy auction as introduced by Ausubel and Milgrom (2002). At each inflection point, the behavior of the ascending proxy auction changes; at a particular payment vector, a bidder will no longer find it profitable to compete for the same bundle or bundles he has been pursuing for the previous rounds and will change his attention to compete for a different bundle. Wurman et al. (2004) provide a mixed-integer linear program (MIP) which traces the payment trajectory, telling at each stage which bundles a bidder is continuing to pursue between inflection points, and the allocation(s) supporting each payment vector along the trajectory. Although this research remains interesting for its ability to “jump” to the points of interest in the ascending proxy auction, we argue that for large-scale practical applications, the trajectory information is irrelevant to the final outcome. Because this additional information (illustrating how the final payments are arrived at through an ascending procedure) is achieved only through additional computational complexity, we argue that the core constraint generation procedure presented here provides

a more viable computational solution for applications in which intermediate payment adjustments can be ignored.

To demonstrate a direct comparison of our algorithm to the inflection point method, we apply our technique on the following worked example of Wurman et al. (2004):

| | A | B | AB | C | AC | BC | ABC |
|---------|----|---|----|---|----|----|-----|
| Buyer 1 | 10 | 3 | 18 | 2 | 18 | 10 | 20 |
| Buyer 2 | 4 | 9 | 15 | 3 | 12 | 18 | 20 |
| Buyer 3 | 1 | 3 | 11 | 9 | 16 | 17 | 25 |
| Buyer 4 | 7 | 7 | 16 | 7 | 16 | 16 | 20 |

This auction has two efficient solutions (both recognized by the inflection point method) with the following outcomes (with payments computed by the inflection point method): Buyers 1, 2, and 3 get A, B, and C, respectively, with $\pi_1 = \pi_2 = 8$, $\pi_3 = 9$; buyer 1 gets A, buyer 2 gets B and C, with $\pi_1 = 8$ and $\pi_2 = 17$. The process solves 11 MIPs to achieve these solutions, each one more computationally complex than the winner-determination problem for the auction.

The core constraint generation algorithm first selects a particular efficient solution, so, for the sake of comparison, we ran our algorithm twice, once for each efficient solution. For the first (with three winners), we solve the initial winner-determination IP, followed by the solution of three winner-determination problems (one with each of the three winners removed) to determine the VCG payments. We then solve a single instance of the separation problem SEP^t , adjust prices once from the VCG payments, and terminate at the solution $\pi_1 = 7.5$, $\pi_2 = 8.5$, $\pi_3 = 9$. In total, five IPs must be solved, each computationally equivalent to the winner-determination problem. For the alternate efficient outcome, we again find the final payments after just a single payment adjustment. For this instance, we solve just four IPs to arrive at the final payments, $\pi_1 = 7.5$, $\pi_2 = 17.5$, with one less IP because there is one less VCG payment to compute.

We first observe that although the total number of IPs/MIPs to find *all* solutions is similar for this example (9 versus 11), in practice only a single auction outcome is necessary, so that the comparison becomes 4 or 5 versus 11. In addition, each instance of the inflection point MIP effectively finds every solution to the winner-determination problem for a particular vector of information released to the proxy, making it necessarily more complex than each IP solved by our algorithm, which finds a single solution to the winner-determination problem.

As the example of Wurman et al. (2004) shows, the ascending proxy and inflection point methods do not necessarily compute the VCG payments, and

may therefore produce a solution with an inferior distribution of deviation incentives.⁶ More importantly though, the added complexity of the MIP introduced by Wurman et al. (2004) is overkill; *every* optimal solution to an \mathcal{NP} -hard winner-determination problem is found for several payment vectors. (The theory of $\#\mathcal{P}$ complexity tells us that finding all solutions of a problem in \mathcal{NP} is fundamentally harder than finding single solutions.) Further, many of the iterations of the procedure reveal information that is irrelevant to the final auction outcome, while each iteration of the core constraint generation algorithm seems to produce (in all investigated instances) a core constraint that is tight (has zero slack) in the final payment outcome.

5.5. Variations on the Core Constraint Generation Algorithm

At no point in the development of the core constraint generation algorithm was it necessary for the separation procedure to begin at the VCG payment vector. In truth, the separation problem SEP^t will find the most violated coalition of bidders starting with an efficient allocation and *any* payment vector. This provides the possibility of many variations of the algorithm, each starting from a different set of payments. Theorem 4 from Ausubel and Milgrom (2002) suggests that the VCG payments should be a good starting point for the algorithm because it will also be the stopping point if the VCG payments are in the core. Additionally, motivated by the approach of Hoffman et al. (2006), and by the beneficial incentive properties suggested by Theorem 3.2 and Corollary 3.3, we propose that the VCG payments provide the strongest starting point for the separation portion of the algorithm, but are open to the possibility of an alternative starting point.

Is it possible that by starting at a different set of initial payments that we reduce the overall computational burden as measured by the number of IPs solved? If we save computational time by not computing the VCG payments, will total computational burden be reduced? The following example demonstrates that the answer may be ambiguous, even for a single auction.

The most natural alternative starting point to consider is the zero payment vector, and we tested this variation of the core constraint generation algorithm on the inflection point example from Wurman et al. (2004). Because the VCG payments are not available, we must alter the formulation $EBPO^t$, giving it some other basis for selecting among all possible payment vectors which minimize the sum of payments within

⁶ One can easily check via Theorem 3.2 that bidder 1 can benefit by unilateral deviation from truth-telling up to one unit under either Wurman solution, while our solution spreads this incentive out, giving 0.5 units of deviation incentive to both bidders 1 and 2.

Table 3 Computational Results Using the CATS Data

| Goods | Bids | WD CPU time | | Total CPU time | | No. of SEP instances | | Total no. of IPs | |
|-------|-------|-------------|----------|----------------|------------|----------------------|---------|------------------|---------|
| | | Worst case | Average | Worst case | Average | Worst case | Average | Worst case | Average |
| 16 | 10 | 0.015 | 0.0018 | 0.015 | 0.0027 | 1 | 0.04 | 8 | 2.96 |
| 16 | 25 | 0.016 | 0.0022 | 0.016 | 0.0102 | 4 | 0.54 | 16 | 4.48 |
| 16 | 50 | 0.016 | 0.0048 | 0.047 | 0.0269 | 6 | 1.22 | 22 | 6.02 |
| 16 | 100 | 0.047 | 0.0103 | 0.359 | 0.1014 | 6 | 1.74 | 22 | 7.76 |
| 16 | 250 | 0.063 | 0.0191 | 0.750 | 0.2190 | 5 | 1.86 | 14 | 7.80 |
| 16 | 500 | 0.297 | 0.0535 | 2.625 | 0.6820 | 10 | 2.68 | 27 | 11.32 |
| 16 | 1,000 | 0.156 | 0.0821 | 2.031 | 1.0610 | 8 | 2.64 | 25 | 11.52 |
| 32 | 10 | 0.015 | 0.0021 | 0.016 | 0.0030 | 2 | 0.06 | 7 | 2.66 |
| 32 | 25 | 0.016 | 0.0028 | 0.032 | 0.0102 | 2 | 0.58 | 10 | 4.68 |
| 32 | 50 | 0.047 | 0.0067 | 0.422 | 0.0395 | 5 | 1.12 | 16 | 6.20 |
| 32 | 100 | 0.172 | 0.0253 | 1.484 | 0.2157 | 5 | 1.56 | 16 | 8.24 |
| 32 | 250 | 0.719 | 0.1432 | 12.032 | 2.5161 | 10 | 3.82 | 26 | 13.86 |
| 32 | 500 | 1.609 | 0.3726 | 39.328 | 9.0243 | 10 | 4.46 | 27 | 15.92 |
| 32 | 1,000 | 1.375 | 0.1832 | 27.422 | 4.8200 | 12 | 5.68 | 32 | 18.66 |
| 64 | 10 | 0.016 | 0.0016 | 0.016 | 0.0031 | 1 | 0.02 | 7 | 2.90 |
| 64 | 25 | 0.031 | 0.0047 | 0.125 | 0.0165 | 7 | 0.88 | 14 | 4.76 |
| 64 | 50 | 0.172 | 0.0226 | 1.438 | 0.2205 | 8 | 1.56 | 16 | 6.64 |
| 64 | 100 | 0.703 | 0.0872 | 11.172 | 1.1957 | 10 | 1.66 | 25 | 9.56 |
| 64 | 250 | 7.516 | 1.8669 | 148.547 | 39.7604 | 31 | 5.14 | 58 | 17.82 |
| 64 | 500 | 48.765 | 5.795 | 2,202.156 | 244.6347 | 33 | 8.92 | 62 | 27.24 |
| 64 | 1,000 | 67.954 | 7.8513 | 3,703.391 | 511.0841 | 22 | 7.82 | 51 | 26.10 |
| 128 | 10 | 0.016 | 0.0012 | 0.016 | 0.00184 | 0 | 0 | 6 | 3.08 |
| 128 | 25 | 0.016 | 0.0029 | 0.094 | 0.0165 | 4 | 0.96 | 12 | 5.54 |
| 128 | 50 | 0.188 | 0.0260 | 2.391 | 0.3298 | 9 | 3.02 | 22 | 10.36 |
| 128 | 100 | 0.938 | 0.1614 | 19.000 | 2.50408 | 10 | 3.1 | 27 | 11.98 |
| 128 | 250 | 58.000 | 6.2901 | 814.235 | 164.5285 | 15 | 5.94 | 44 | 23.98 |
| 128 | 500 | 1,381.656 | 155.0251 | 51,824.359 | 6,812.6973 | 169 | 20.26 | 218 | 47.10 |

Note. Each data point represents 50 simulated auction instances. All times are in seconds. The total number of IPs represents SEP instances together with the initial winner determination instance and one IP to compute each winning bidder's VCG payment.

the core. Although there are several possibilities, we used the most simple: simply replace the VCG payments in the EBPO^t formulation with zeros, so that the algorithm finds a payment vector that minimizes the maximum payment of any bidder among all payment vectors that minimize total payments.

We ran this alternative algorithm twice on the example from Wurman et al. (2004), again, once for each efficient solution. For the first efficient solution (with three winners), we solve six IPs: the initial winner-determination problem, and five instances of the separation problem. Compared with the five total IPs solved when using the VCG starting point, this technique seems inferior. Surprisingly however, if we choose the alternative efficient solution for the same problem (with just two winners), we solve only three IPs, compared with four when starting at the VCG payments. Clearly, there is no dominance for one technique over the other in terms of computational burden alone, even for the same auction! (The results are identical in both final payments and number of IPs solved if we use the equal-pay secondary objective discussed in §5.2.)

Interestingly though, starting at payments of zero and minimizing the maximum difference from zero

payment over all bidders, we find exactly the same payments as we do using the inflection point method (and hence the generic ascending proxy auction) for this example. As noted earlier, this solution has inferior incentive properties, causing us again to lean toward the use of a VCG starting point. The less costly starting point does not always lead to overall improvement of performance and forfeits the ability to approximate VCG payments explicitly.

5.6. Some Computational Experiments

Table 3 summarizes our experience implementing instances of the core constraint generation algorithm to solve instances generated using CATS (Combinatorial Auction Test Suite) to model real-life auctions. Leyton-Brown et al. (2000) provide details of how combinatorial auction XOR bids are generated to emulate several economically relevant auction profiles using CATS. In our case, we generated 50 auction instances for several arbitrarily chosen combinations of goods and bids. Using the “arbitrary hybrid” feature of CATS, we randomly selected a distribution from among the several economic scenarios devised by CATS with equal probability for each instance. (We omit the legacy distributions which are not intended,

as the rest of CATS is, to reflect real-life implementations.) All other parameters were left at their default values for CATS 2.0. Computations were performed using CPLEX 9.0 on a 3.2 GHz Pentium 4 CPU with 1.0 GB of RAM.

As indicated by this data, bidder-Pareto-optimal core payments can be computed using the core constraint generation technique at a reasonable pace for many realistically sized auctions. With even the largest instance taking around 14 hours, our experience with these 1,350 instances indicates that it would be reasonable to expect that the outcome of a sealed-bid combinatorial auction in this size range could be reported by the auctioneer on the following day, if not in the first few seconds for smaller instances.

6. Concluding Remarks

Recent trends indicate that combinatorial auctions, in which bidders submit bidding information on many bundles, are fast becoming viable for several large-scale business-to-business and governmental auctions. Because of their adherence to efficient allocations, core outcomes which cannot be challenged by disgruntled bidders, and bidder-Pareto-optimality, guaranteeing that bidders do not suffer high payments unnecessarily, bidder-Pareto-optimal mechanisms are especially attractive solutions for the allocation of public goods. Ausubel et al. (2006) provide the practical clock-proxy-auction design, motivated by a great deal of real experience conducting high-stakes combinatorial auctions. Further, this auction format has received a great deal of attention as a candidate for existing and proposed governmental auction applications (Kwerel 2004). Consequently, the improvements described in this paper, addressing critical issues related to the proxy auction and bidder-Pareto-optimal mechanisms in general, promise to have a wide-ranging impact.

The clock-proxy auction incorporates a demand revelation phase which is followed by (and terminates with) a proxy auction, which is essentially a sealed-bid combinatorial auction with a particular payment rule. This sealed-bid phase can be run a number of different ways, but as laid out in the earlier ascending proxy auction of Ausubel and Milgrom (2002), there are only a few essential features: efficiency based on reported values, and the use of a payment rule that finds a bidder-Pareto-optimal (as opposed to a VCG or pay-as-bid) point in the core. In this paper, we have presented a new “direct” method of finding auction outcomes which meet these same criteria, promising to deliver faster results in the practical environment, to allow for auctions with a greater number of items (relevant, for example, to the large number of landing-slots in proposed FAA applications), and to assure bidders of a transparent paradigm for fair payments.

Along the way, we refined several key concepts governing these auction outcomes. A few general results indicated the virtue of the bidder-Pareto-optimal mechanism, justifying an approach which minimizes total payments within the core primarily, and minimizes the maximum difference from VCG payments as a secondary objective. Our treatment provides a more clear understanding on the selection of a bidder-Pareto-optimal point from the core, where in the past, this notion was somewhat vague.

As the primary justification for total-payments minimization within the core, Propositions 5.1 and 5.2 describe a form of collusion that is nullified using total payment minimization, while the example of Figure 1 demonstrates how bidders can collude to achieve a “socially unacceptable” outcome using alternative techniques. As a secondary benefit of total-payment minimization, Corollary 3.3 demonstrates that total individual incentives to deviate are also minimized. As discussed in §3, this property gains importance in the presence of costly information.

As a final remark on the selection of a core outcome, we note that the public perception of fairness may be the strongest motivation to adopt a total-payment-minimization mechanism. For governmental auctions in particular, many will take comfort that the auctioneer is not selecting an outcome according to self-interest, but in the combined best interests of the bidders. Indeed, the Air Transportation Association, representing the concerns of airlines who would bid in the recently proposed auctions for airport landing slots, voices a concern that “slot auctions should not be designed to maximize payments from airlines” (*Airline Business Report* 2005, p. 1). What could be a more accommodating response to these public demands than a mechanism assuring total-payment minimization over the space of socially acceptable outcomes? In private sector B2B auctions, on the other hand, assurance of an auction that reaps benefits in a minimal fashion may be the best method to maintain long-term trade relationships and encourage repeated participation. In both cases, only the core constraints, representing competition, should drive up the payments, not the artificial device used to determine payments. Looking back at the classical auction literature, we see that the second-price single-item auction also minimizes payments with respect to core constraints, reinforcing our payment mechanism as a viable generalization of the second-price auction to the realm of multiple goods.

Additionally, the core constraint generation method offers computational benefits over other techniques described in the recent literature. By solving the separation problem at each point in payment space, we only consider coalitions that actually threaten to block a potential outcome, obviating the need to apply the

entire exponential set of constraints. Our definition of core constraints (as in CORE) are equivalent to those defined as core constraints in Hoffman et al. (2006) and Parkes (2002) through a linear change of variables, but reformulating the constraints in terms of payments, rather than payoffs, provides the proper intuition for our observations on coalitional contribution, facilitating the formulation of the separation problem. To our knowledge, the formulation of the separation problem for violated core constraints and the development of a price adjustment procedure utilizing this approach are novel contributions to the literature.

Further, the core constraint generation algorithm takes place without the need for a price increment as in the ascending proxy technique, accelerated or otherwise. EBPO^t upholds every generated core constraint for the remainder of the procedure. This ensures that the same coalition will not appear repeatedly as blocking in our procedure. Experience with the ascending proxy method, on the other hand, shows that after an increment and re-resolution of the winner-determination problem, the same coalition may appear for several iterations.

Wurman et al. (2004) find a different way around this issue with the use of a MIP that jumps to each “change point” or “inflection point” at which the relevant coalitions change, but we argued against that technique for several reasons. Most importantly, the inflection point method finds more information than is needed for most practical implementations, at the cost of significant additional computational complexity.

The strength of the ascending proxy technique (with the accelerating techniques of Hoffman et al. 2006, the inflection point method of Wurman et al. 2004, or otherwise) is that an auction with iterative revelation of demand is either performed explicitly or simulated. But, because the information is submitted to a proxy (and cannot be changed by the bidder) iterative revelation of demand is strategically irrelevant, and merely makes a nice story to explain the process to bidders. We propose the core constraint generation procedure as an alternative that bypasses the iterative revelation of demand to achieve the final outcome more directly. Only allocations that affect the final payments of the winning bidders are considered; many allocations that are only feasible at intermediate, irrelevant payment vectors can be ignored.

Our core constraint generation algorithm shows much promise as the method of choice for the determination of bidder-Pareto-optimal core outcomes in public sector combinatorial auctions. As the examples provided in this paper show, our technique clearly dominates the proxy-auction algorithm, the accelerated proxy developed by Hoffman et al. (2006), and

the inflection point approach of Wurman et al. (2004). We hope that ongoing studies (performed by the authors and by researchers associated with the FAA and FCC) testing large-scale implementations will confirm the rapidness, robustness, and scalability of our approach, thereby providing a practical methodology for computing fair solutions in public sector combinatorial auctions.

7. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at <http://mansci.journal.informs.org/>.

Acknowledgments

The authors gratefully acknowledge summer research support from their respective institutions, as well as the National Science Foundation (award number DMI-0205489).

References

- Airline Business Report. 2005. Gulf between legacy carriers, auction proponents revealed by mock proposal. *Airline Bus. Rep.* 23(2).
- Ausubel, L. 2006. An efficient dynamic auction for heterogeneous commodities. *Amer. Econom. Rev.* 26(3) 602–629.
- Ausubel, L., P. Cramton. 1999. The optimality of being efficient. Working paper, University of Maryland, College Park, MD.
- Ausubel, L., P. Milgrom. 2002. Ascending auctions with package bidding. *Frontiers Theoret. Econom.* 1(1) 1–42.
- Ausubel, L., P. Cramton, P. Milgrom. 2006. The clock-proxy auction: A practical combinatorial auction design. P. Cramton, Y. Shoham, R. Steinberg, eds. *Combinatorial Auctions*, Chap. 5. MIT Press, Cambridge, MA, 115–138.
- Bazaraa, M., H. Sherali, C. Shetty. 1979. *Nonlinear Programming: Theory and Algorithms*. John Wiley and Sons, New York.
- Clarke, E. 1971. Multipart pricing of public goods. *Public Choice* 11(1) 17–33.
- de Vries, S., J. Schummer, R. V. Vohra. 2007. On ascending Vickrey auctions for heterogeneous objects. *J. Econom. Theory* 32(1) 95–118.
- Grötschel, M., L. Lovasz, A. Schrijver. 1981. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica* 1(2) 169–197.
- Groves, T. 1973. Incentives in teams. *Econometrica* 41(4) 617–631.
- Günlük, O., L. Ladányi, S. de Vries. 2005. A branch-and-price algorithm and new test problems for spectrum auctions. *Management Sci.* 51(3) 391–406.
- Hoffman, K., D. Menon, S. van den Heever, T. Wilson. 2006. Observations and near-direct implementations of the ascending proxy auction. P. Cramton, Y. Shoham, R. Steinberg, eds. *Combinatorial Auctions*, Chap. 17. MIT Press, Cambridge, MA, 415–450.
- Kwerel, E. 2004. Personal communication. NEXTOR Workshop: Government, the Airline Industry and the Flying Public: A New Way of Doing Business.
- Leyton-Brown, K., M. Pearson, Y. Shoham. 2000. Towards a universal test suite for combinatorial auction algorithms. *EC '00: Proc. 2nd ACM Conf. Electronic Commerce*. ACM Press, New York, 66–76.
- Nash, J. 1950. The bargaining problem. *Econometrica* 18(2) 155–162.

- Parkes, D. 2001. Iterative combinatorial auctions: Achieving economic and computational efficiency. Ph.D. thesis, Computer and Information Science, University of Pennsylvania, Philadelphia, PA.
- Parkes, D. 2002. Notes on indirect and direct implementations of core outcomes in combinatorial auctions. Technical report, Harvard University, Cambridge, MA.
- Parkes, D., J. Kalagnanam, M. Esó. 2001. Achieving budget-balance with Vickrey-based payment schemes in exchanges. *Proc. 17th Internat. Joint Conf. Artificial Intelligence (IJCAI'01), Seattle, WA*, 1161–1168.
- Rothkopf, M., R. Harstad. 1995. Two models of bid-taker cheating in Vickrey auctions. *J. Bus.* **68**(2) 257–267.
- Rothkopf, M., A. Pekeč, R. Harstad. 1998. Computationally manageable combinatorial auctions. *Management Sci.* **44**(8) 1131–1147.
- Rothkopf, M., T. Teisberg, E. Kahn. 1990. Why are Vickrey auctions rare? *J. Political Econom.* **98**(1) 94–109.
- Sakurai, Y., M. Yokoo, S. Matsubara. 2000. An efficient approximate algorithm for winner determination in combinatorial auctions. *EC '00: Proc. 2nd ACM Conf. Electronic Commerce*. ACM Press, New York, 30–37.
- Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *J. Finance* **16**(1) 8–37.
- Wurman, P., J. Zhong, G. Cai. 2004. Computing price trajectories in combinatorial auctions with proxy bidding. *Electronic Commerce Res. Appl.* **3**(4) 329–340.

e - companion

ONLY AVAILABLE IN ELECTRONIC FORM

Electronic Companion—“Fair Payments for Efficient Allocations in Public Sector Combinatorial Auctions” by Robert W. Day and S. Raghavan, *Management Science*, DOI 10.1287/mnsc.1060.0662.

Online Appendix

Proof of Theorem 3.1

First, we prove the initial statement, that any bidder-Pareto-optimal core point corresponds to a Nash equilibrium in semisincere strategies when using a mechanism that is bidder-Pareto-optimal with respect to submitted preferences. To construct a proof by contradiction, suppose that for any specific point in the core (π^A, A) , all bidders bidding $b_j(\cdot) = \max(v_j(\cdot) - v_j(S_j^A) + \pi_j^A, 0)$ is not a Nash equilibrium. Thus, there is a bidder l who may profitably deviate from this bid profile (unilaterally), achieving a new winning coalition D with payment vector π^D , and allocating bundle S_j^D to each bidder j . It may be worth noting that l must have been a winner in the allocation A . (If not, l must bid more than b_l on some bundle, which implicitly equals v_l , because when losing $S_l^A = \emptyset$ and $\pi_l = 0$. But when this losing bidder increases a bundle bid to the minimum amount necessary to be included in an efficiently determined allocation, she will pay her bid, which would be greater than her value, contradicting the profitability of the deviation.)

The profitability of the new deviation strategy for bidder l implies that $v_l(S_l^A) - \pi_l^A < v_l(S_l^D) - \pi_l^D$, or equivalently, $v_l(S_l^D) - v_l(S_l^A) + \pi_l^A > \pi_l^D$. Because the left side of this inequality is exactly the bid of l prior to deviation, we may write this as follows:

OBSERVATION 1. $b_l(S_l^D) > \pi_l^D$.

Next, note that because of bidder-Pareto optimality in the core of the mechanism (with respect to submitted bids) and the bid profiles $b_j(\cdot)$, for any winning bidder j there is a coalition not including j whose maximum total bids on the entire set of auction items is equal to the payments of the winners. If this were not true, the mechanism could find a Pareto improvement, lowering the payment of j without disturbing the surplus of other players and without leaving the (apparent) core. In particular, for the initial allocation to the coalition A , we note the following:

OBSERVATION 2. There is a coalition B such that $l \notin B$ and $\sum_{j \in B} b_j(S_j^B) = \sum_{j \in A} \pi_j^A$.

Thirdly, we note that under the initial strategies (prior to deviation by l) we have $\sum_{j \in A} \pi_j^A \geq \sum_{j \in D} (v_j(S_j^D) - v_j(S_j^A) + \pi_j^A) = \sum_{j \in D} b_j(S_j^D)$, with the first inequality following because (π^A, A) is in the core with respect to true valuations. We re-state this as the following:

OBSERVATION 3. $\sum_{j \in A} \pi_j^A \geq \sum_{j \in D} b_j(S_j^D)$.

Putting these facts together, and considering what bidders in B are willing to offer the auctioneer after the deviation of bidder l according to their own undeviated bids, we have the following string of inequalities:

$$\begin{aligned}
 \sum_{j \in B} (b_j(S_j^B) - b_j(S_j^D) + \pi_j^D) &= \sum_{j \in A} \pi_j^A + \sum_{j \in B} (\pi_j^D - b_j(S_j^D)) \quad \text{by Observation 2 and the selection of } B \\
 &\geq \sum_{j \in D} b_j(S_j^D) + \sum_{j \in B} (\pi_j^D - b_j(S_j^D)) \quad \text{by Observation 3} \\
 &> \pi_l^D + \sum_{j \in D \setminus l} b_j(S_j^D) + \sum_{j \in B} (\pi_j^D - b_j(S_j^D)) \quad \text{by Observation 1}
 \end{aligned}$$

$$\begin{aligned}
&= \sum_{j \in D} \pi_j^D + \sum_{j \in B \setminus D} (\pi_j^D - b_j(S_j^D)) \quad \text{by canceling terms in } B \cap D \\
&= \sum_{j \in D} \pi_j^D
\end{aligned}$$

because each $S_j^D = \emptyset$ for bidders not in D , and thus $\pi_j^D = b_j(S_j^D) = 0$ for all bidders j in $B \setminus D$. Altogether this implies that

$$\sum_{j \in B} (b_j(S_j^B) - b(S_j^D) + \pi_j^D) > \sum_{j \in D} \pi_j^D.$$

But with Lemma 4.1, this implies that the coalition B would be willing to pay strictly more for the bundles S_j^B following the deviation by bidder l than coalition D actually pays following the deviation, contradicting the core property of the mechanism based on the submitted bids. Thus, every bidder-Pareto-optimal core point has a semisincere bid profile supporting that point as a Nash equilibrium.

To prove the second claim in the statement of the theorem, suppose that we have strategies $b_j(\cdot)$ for each j that together form a Nash equilibrium with winning coalition A who pay π_j^A and each receive bundle S_j^A , and that this outcome is not in the core with respect to true valuations. That is, there is some coalition B that objects to the the payments π_j^A :

$$\sum_{j \in A} \pi_j^A < \sum_{j \in B} (v_j(S_j^B) - v_j(S_j^A) + \pi_j^A)$$

even though

$$\sum_{j \in A} \pi_j^A \geq \sum_{j \in B} (b_j(S_j^B) - b_j(S_j^A) + \pi_j^A)$$

by the core property of the mechanism with respect to submitted bids. Rearranging we have $\sum_{j \in B} (v_j(S_j^B) - b_j(S_j^B)) > \sum_{j \in B} (v_j(S_j^A) - b_j(S_j^B))$ and thus $v_l(S_l^B) - b_l(S_l^B) > v_l(S_l^A) - b_l(S_l^A)$ for at least one bidder $l \in B$. But this says that in a Nash equilibrium bidder l shades (reduces her bid) more on bundle S_l^B than on bundle S_l^A , while Theorem 3a of Ausubel and Milgrom (2002) shows that a semisincere strategy with equal shading on all bundles is always a best response for every bidder when the mechanism is bidder-Pareto optimal in the core with respect to the submitted bids, a contradiction. (Theorem 3a of Ausubel and Milgrom 2002 is stated in terms of the ascending proxy auction, but because the only property of that auction used in the proof is bidder-Pareto optimality in the core with respect to the submitted bids, the result immediately generalizes.) Thus, every Nash equilibrium profile of bids results in a core outcome with respect to true valuations.

Finally, we must show that every Nash equilibrium in semisincere strategies is bidder-Pareto optimal within the core. If a collection of π_j^A s and S_j^A s represent a Nash equilibrium that is not bidder-Pareto optimal in the core, then there is some alternative bidder-Pareto-optimal outcome in the core with π_j^B and S_j^B for each bidder and at least one bidder l that prefers outcome B , while every other bidder is at worst indifferent. By deviating to $b_l(\cdot) = \max(v_l(\cdot) - v_l(S_l^B) + \pi_l^B, 0)$ bidder l can obtain utility of $v_l(S_l^B) - \pi_l^B$ because no other outcome is bidder-Pareto optimal in the core with respect to the preferences reported following the deviation. By our selection of l this deviation obtains positive gains, contradicting our assumption of a Nash equilibrium. \square

Reference

See references list in the main paper.