Pricing Valid Cuts for Price-Match Equilibria

Robert Day
Department of Operations and Information Management, University of Connecticut, Storrs, Connecticut 06269
robert.day@uconn.edu

Benjamin Lubin
Department of Information Systems, Boston University, Boston, MA 02215
blubin@bu.edu

We use valid inequalities (cuts) of the binary integer program for winner determination in a combinatorial auction (CA) as "artificial items" that can be interpreted intuitively and priced to generate Artificial Walrasian Equilibria. While the lack of an integer programming gap is sufficient to guarantee a Walrasian equilibrium, we show that it does not guarantee a "price-match equilibrium" (PME), a refinement that we introduce, in which prices are justified by an iso-revenue outcome for any hypothetical removal of a single bidder. We prove the existence of PME for any CA and characterize their economic properties and computation. We implement *minimally artificial* PME rules and compare them with other prominent CA payment rules in the literature.

Key words: combinatorial auction, Walrasian Equilibrium, prices, cut generation

1. Introduction

Walrasian Equilibrium (WE) prices support an efficient market outcome in which each bidder weakly prefers the bundle she is awarded by the market at the provided prices, and in which an item receives a positive price only if there is no excess supply of the item. A Combinatorial Auction (CA) is a market in which bidders express detailed preference information over combinations of indivisible items. In spite of the importance of WE as a foundation of microeconomic theory, it is well known that WE item prices may not exist in a general CA.

Bikhchandani and Mamer (1997) first showed in the CA context that a WE exists if and only if the linear programming relaxation of its *winner determination problem* (WDP-LP) has an integer optimal solution. Here, we emphasize that their result holds in the context of linear anonymous prices on the *original items* of interest, i.e., the items listed by the auctioneer for sale. Following the spirit of this approach, we too apply strong duality in the WDP-LP context to obtain WE prices, but we also generate valid inequalities or cuts to the WDP and then price those cuts to establish WE prices in any case where they would not otherwise exist. We call these *Artificial Walrasian Equilibria* or AWEs. Philosophically, these cuts offer prices for *artificial items* that the auctioneer "didn't think to offer" for sale, such as a license to own a package of a particular size, or for example, a license to take any pair of items from the set $\{A, B, C, D, E\}$. If the auctioneer

had the foresight to set up and sell such licenses, then a normal WE would be available. Here, we let the auction generate these artificial items endogenously, sparing the auctioneer from the need to predetermine which combinatorial structure of artificial items would do the job.

Recall that the core of a game is the set of payoffs that cannot be *blocked* by a coalition that could do better by trading among themselves. We show that there are many ways to meaningfully generate and price cuts as artificial items for a given CA, proving that any point in the core of the auction game can be supported as an AWE, and that a single core outcome may have several representations as an AWE. Given this large space of AWEs to explore, we introduce an important subset of WE prices which we call *price-match equilibrium* (PME) prices, defined here as item prices that would still be supported if any one bidder were removed (without charging any other bidder more than her bid). Such prices are more *justified* from each bidder's perspective than non-PME WE prices: each player observes that should they leave, the seller could still collect any price they refuse to pay.

We show that PME prices always exist for a CA when artificial items are permitted. While not previously described in the literature, we emphasize the PME concept as a meaningful normative criterion for market design: any non-PME price outcome should be viewed as an "unreasonable" over-pricing by the seller, an artifact of applying the WE concept to general combinatorial markets naïvely, and a practice that should be avoided. Each bidder should demand that the auction not charge her more than can be *matched* by other bidders. In procurement auctions, such as those in public contracting, the direction reverses: working contractors should demand not to be underpaid simply because the auctioneer selected a non-PME linearization of the combinatorial market. Every bidder should demand that the amounts they pay (or are paid in procurement) are fully matched by a feasible set of competing bids. Through artificial-item pricing as presented here, this refined notion of price-matching (a stricter notion of fairness than WE) is available.

We show that minimum-revenue core-selecting (MRC) payments (as introduced in Day and Raghavan 2007) can always be supported with PME prices, but other PME possibilities are also available. As a new alternative to MRC prices, we refine further to suggest the use of *Minimally Artificial* PME payments, or MAP payments for short. These are selected such that the aggregate adjustment associated with the use of cuts is minimized, subject to PME being established, a way to stay "as close to standard linear prices as possible." Generic MAP payments are then *elaborated upon* to find a minimal set of specific supporting cuts (and thus specific linear prices) that uphold the MAP property.

Final MAP prices will sometimes correspond to existing payment schemes (like MRC, or VCG), but given other bids they may be distinct from any fixed pricing approach from the literature that we are aware of. The total-revenue characterization for a fixed set of bids is given by:

$$\text{NWE} \geq \text{MAP} \geq \text{MRC} \geq \text{VCG} \tag{1}$$

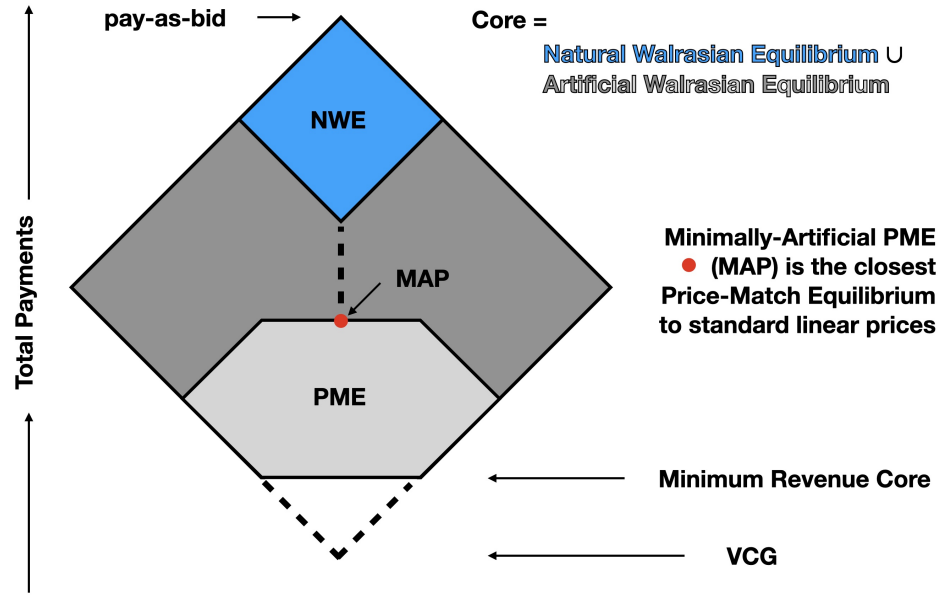


Figure 1 **A generic diagram of the core in payment space with five alternative CA pricing rules, including our newly introduced MAP prices and the PME region of prices. The central dark gray region consists of non-PME points that can be supported by AWE prices.**

where NWE represents any *Natural* WE (i.e., without any cuts/priced artificial items) should they exist, and VCG is the classical *Vickrey-Clarke-Groves* mechanism, known to be the unique dominant-strategy incentive-compatible mechanism under typical assumptions. We show through a series of examples that each inequality above can be strict in some CA instances and exact equality in other instances. When no NWE exists, we may still have non-PME AWEs with strictly greater revenue than MAP prices, or at the other extreme, for example when bidders have unit-demand, minimal NWE = VCG and the distinction among these four solution concepts disappears. We further show that unlike MRC payments, MAP prices do not suffer from the *unrelated-goods problem* introduced by Beck (2013), and that the commonly studied gross substitutes condition, while sufficient to imply NWE existence, does not imply that the NWE is a PME.

Overall, MAP prices emerge as an alternative pricing paradigm for both prescribing total payments and *explaining* them, precisely where traditional WE fails and where MRC and VCG had previously offered only opaque lump-sum payments. In some cases, the latter two revenue inequalities may be equality, and so MRC and VCG may admit supporting MAP prices, which therefore offer an explainable breakdown into individual linear prices. In other cases, MAP prices will differ from VCG and MRC.

Figure 1 summarizes and contextualizes our contributions with a generic depiction of the core in payment space. The core itself is bounded above by bids (individual rationality constraints) and below by $\geq$-constraints that ensure subsets of winning bidders must weakly exceed the offers of

potentially blocking coalitions. NWE are defined by envy-freeness constraints: the prices of the items I bid on must be large enough that I don't want to switch to a different package. The blue NWE region is known to shrink to non-existence, or expand to fill the entire core, including perhaps the VCG point, depending on the structure and strength of the bids. The region of PMEs, by contrast, are defined by $\leq$-constraints: payments must not be so much that they cannot be matched by losing bids. The PME region may contain outcomes with greater total payments than MRC (as depicted) or may shrink to coincide with the MRC exactly, but PME always includes the MRC points, which themselves must always exist.

Pay-as-bid and the VCG payments represent the two extreme benchmark pricing rules for CAs in Figure 1, outcomes that have been broadly studied and need little introduction. The latter are known to sometimes be outside the core, as depicted, a motivation for the use of MRC points which are by definition the smallest total-payment outcomes in the core. Together, VCG, MRC, NWE (when they exist), and pay-as-bid, give four potential payment outcomes for a given set of bids, all of which have been studied extensively and used in practice. Here, we introduce and explore the use of MAP prices as a fifth well-defined alternative outcome for any CA, an intermediate point which, as shown, often lies in the interior of the core.

To see one potential application of the results presented here, suppose MAP was used in the final (supplementary) round of the widely used combinatorial clock auction (Ausubel et al. 2006). If the market were to clear (with no excess supply) in the clock phase of such a CA, and if the associated NWE satisfies the newly defined PME property, a MAP auction would terminate at NWE with no change in payments between the clock phase and final outcome, given straightforward bidding. This would avoid the potentially hard-to-explain drop in revenue a seller could experience under a VCG or MRC final pricing rule, when agreeable NWE prices were already obtained in the clock phase. Example 4 below shows an example of this type; clock prices would find NWE = MAP payments of 20 to clear the market, but a supplementary round with VCG or MRC would lower revenue to 10. Importantly, in other situations (like Example 7 below) all NWE may fail the PME condition, while a MAP price rule for the supplementary round would produce the more explainable PME outcome. Given the above revenue characterization, MAP prices offer an intermediate payment scheme for the market-design tool kit, between the traditional schemes in both revenue generation and truthfulness, which tend to be directly opposed as described by Parkes, Kalagnanam, and Eso (2001).

2. Background and Related Literature

Real-world CA applications often use a combinatorial-clock format, with a first stage consisting of iterative "clock" rounds, price-vector announcements eliciting demand responses, followed by

a sealed-bid CA final stage, as in Ausubel, Cramton, and Milgrom (2006). Here, we focus on a the latter sealed-bid CA as a standalone process, which can be used as presented or as the final sealed-bid phase of a more complex mechanism. For such a simple mechanism, bidders submit collections of bids and the auctioneer announces final winners and payments on behalf of the seller. Alternatively, this paper studies a *solution concept* for the auction game, interesting in its own right as an explainable, agreeable outcome when all values are known.

We assume that we are working on problems of a practical size, ones for which we can comfortably solve the NP-hard WDP with available software, such as CPLEX or Gurobi. Under similar assumptions, in 2013 the UK used CAs to sell 28 licenses grouped into 4 types, then in 2014 Canada sold 98 licenses in 56 types, and then in 2015 318 licenses of 106 types, the latter aided by the use of a specialized *OR*-based bidding language (Ausubel and Baranov 2017). Thus, while a generic CA winner-determination problem is NP-hard and has poor asymptotic performance as problem size grows, many practical problem instances are far enough from such asymptotic behavior for the limitation to be of little concern.

Walrasian Equilibrium (aka competitive equilibrium) is a central solution concept in microeconomics, with literature far too broad to be surveyed here. (See Bichler et al. 2021, for a recent survey from the optimization perspective.) Among related work, Baldwin et al. (2022) recently provided a strong-substitutes bidding language, a specific submission format (and thus preference restriction) for CAs in which Walrasian equilibrium prices (on the natural items) are guaranteed to exist. Bichler et al. (2018) apply WE principles to a real-world application in the regulation of fisheries, but in their case, when linear anonymous prices do not exist, they relax efficiency in favor of keeping linearity and anonymity of the prices (for natural items) strictly. Here, we instead explore artificial items to approximate WE, relaxing along a different dimension while strictly maintaining efficiency with respect to submitted bids.

Approximation to competitive equilibrium is also used in allocation markets with fictional money by Budish (2011), a format that has proved useful in practice for seats in academic courses and other situations of heterogeneous preferences and high-demand. More recent literature along this thread is given by Nguyen and Vohra (2022) who investigate social-approximate equilibria when no WE exists, providing both approximation techniques via relaxation and preference restrictions that result in provable bounds on the relaxation error. Taken as a whole, this literature stream explores a number of ways to approach practical market design by approximating WE along various dimensions when they are not guaranteed to exist, as in the current work.

When the coalitional valuation function (WDP considered over any subset of bidders) is *submodular* the minimal NWE and VCG outcomes are guaranteed to be equivalent (de Vries et al. 2007), and therefore the revenue comparison (1) above flattens, yielding strict equality among the

NWE, MAP, MRC, and VCG solution concepts. (Coalitional submodularity holds when any bidder always adds weakly more value to a set of bidders than any of its supersets.) Under such a strict assumption on preferences, the auction of that paper can be run to deliver the same outcome we prescribe here, VCG with a decomposition into linear item prices (a WE). Here, we make no such assumption but additionally provide distinct MAP prices in cases where coalitional submodularity fails.

Also related to the current work are others that implicitly introduce cuts and price them. Among these, Milgrom and Watt (2022) find market outcomes that approximate natural WE when none exist. They do not make the distinction of desiring PMEs as introduced here, nor do they explore the concept of cuts as artificial items directly. Like O'Neill et al. (2005), they have a single markup offset which can be interpreted in the language of the current paper as a single cut to the WDP-LP based on prices, as if a tax on auction spending. The combinatorial cuts of the current paper attempt instead to generate artificial structures that are as *simple* or *explainable* as possible, in that we look for cuts reflecting simple relationships among bids that can be represented as integer knapsack constraints, with smaller constant terms favored as simpler and more explainable.

Even more closely related is the recent work of Lahaie and Lubin (2019) who introduce the adaptive (endogenous) generation of price terms in a dynamic (tâtonnement) CA. They focus on CAs with unique items (supply of one item for each commodity) and "monomial" price terms that each correspond to a subset of items, allowing for the pricing of certain synergies directly, with the goal of terminating at WE prices. From the perspective of the current paper, each monomial price term can be viewed as a clique cut (i.e., a fictional item for which there is exactly one copy), though it can also be shown that not all clique cuts are implicitly considered. (For example, Example 1 below uses a clique cut that does not correspond to a unique subset or monomial term in that paper.) They are able to show computationally that these cuts are highly effective in generating WEs by augmenting the space of prices under consideration, complementing the perspectives offered here. Here we also show a framework that adds pricing terms, but we move to the more general multi-unit CAs and explicit search over a broader set of possible cuts.

Similarly, the concurrent work of Emadikhiav and Day (2024) uses the full class of clique cuts (but not cuts with an arbitrary constant term) to extend the price space in an attempt to better approximate WE prices when a NWE does not exist. The focus there is on larger problems and compact representation, rather than the details of finding AWE with desirable properties. Notably, the PME property is not considered in any of these related papers, which to the best of our knowledge is described for the first time here.

All AWE (and hence PME) prices naturally produce *core-selecting* auctions, in which no coalition of bidders could propose a mutually (weakly) preferred outcome with higher total revenue. Core-selecting auctions are also described in Day and Milgrom (2008), with the precise selection of core

points discussed in Day and Cramton (2012) and Bünz et al. (2022). We use the ideas of quadratic minimization as a tie-breaking rule as in the latter papers, but emphasize that here, quadratic tie-breaking is over item prices instead of personal total payments, an arguably better approach. Here we also show that all core-selecting points can be reached as an AWE via just one valid cut, which allows us to easily select the core point we want: PME prices with a minimal amount of artificial price adjustment, decomposed into a collection of more interpretable cuts.

3. Model Preliminaries

In a *combinatorial auction* (CA), a single seller offers J different types of indivisible items for sale to a set of I bidders, denoted $\mathcal{I} = \{1, 2, ...i, ...I\}$. The set of item types is denoted $\mathcal{J} = \{1, 2, ...j, ...J\}$, and each item j has a supply of $c_j \geq 1$ identical copies available to be purchased. CAs with unique items, i.e., those with $c_j = 1$ for all j, are often studied in the literature; here we allow for the more general case. Column vector c will denote the complete supply vector with elements c_j.

Each bidder i is interested in receiving a *bundle* of items, represented as an integer column vector $a \in \mathbb{Z}^J_{\geq 0}$, with $0 \leq a_j \leq c_j$ for each component a_j. Preferences are represented by a *valuation function* $v_i : \mathbb{Z}^J_{\geq 0} \to \mathbb{R}_{\geq 0}$, maintaining standard assumptions of normalization and monotonicity (aka free disposal), i.e., $v_i(\emptyset) = 0$, and $a' \geq a \implies v_i(a') \geq v_i(a)$, for all $i \in I$ and any bundles a' and a.

In a sealed-bid CA, the auctioneer collects a set of bids $\mathcal{K} = \{1, 2, ...k, ...K\}$ from the bidders, each one containing a bundle of interest a^k and a monetary bid amount b_k. The subset of bids made by a particular bidder i is $\mathcal{K}_i$, and when $|\mathcal{K}_i| = 1$ we say bidder i is *single-minded*. Where convenient, we denote as $i(k)$ the unique bidder making bid k. While many "bid languages" have been considered in the literature for particular markets, general studies of CAs typically focus on this XOR language, allowing for full expressivity. Stated simply, this means that the auctioneer may accept at most one bid from $\mathcal{K}_i$ for any bidder i.

The bid amounts b_k are collected into a row vector b, and the indexed set of all bundles (column vectors) associated with bids a^k can be collected into a $J \times K$ matrix $\boldsymbol{A}$, i.e., each element $\boldsymbol{A}_{j,k} = a^k_j$ of $\boldsymbol{A}$ equals the desired quantity of item j in bid k. As in other literature on Walrasian equilibria in CAs, we consider the existence and computation of prices as a solution concept while assuming bidders tell the truth. Thus, except where otherwise noted we assume truthful bidding, i.e., $b_i(a) = v_i(a)$ for all i and a.

For small illustrative examples, we write $b_i(a)$ for bidder i's bid amount on bundle a, or alternatively, for example, $b_3(ABD) = 20$ to denote bidder 3's bid is \$20 for the set $\{A, B, D\}$, with an alternative naming of items by capital letters. We drop set braces in many places where the meaning should be clear, e.g., writing $\mathcal{I} \setminus i$ instead of $\mathcal{I} \setminus \{i\}$, and will say, for example, $b_1(9A) = 85$ for a bid of 85 from bidder 1 for 9 copies of item A.

3.1. Winner Determination

To capture the XOR nature of these bids, let $\boldsymbol{B}$ be the $I \times K$ matrix with elements $\boldsymbol{B}_{i,k} = 1$ when $k \in \mathcal{K}_i$ and $\boldsymbol{B}_{i,k} = 0$ otherwise. The standard winner determination problem (WDP) for the CA can be written as a binary linear program:

$$w(\mathcal{I}, \mathcal{J}, \mathcal{K}) = \quad \max \quad bx \tag{WDP}$$

$$\text{s.t.} \quad \boldsymbol{A}x \leq c \tag{2}$$

$$\boldsymbol{B}x \leq \mathbf{1} \tag{3}$$

$$x \in \{0,1\}^K \tag{4}$$

where $\mathbf{1}$ is the vector of all ones of dimension I. This optimization selects a bid (column) k by setting the corresponding decision variable x_k to 1, with constraints (2) ensuring that no more than c_j copies of any item is sold, and (3) upholding the XOR language, i.e., at most one bid may be accepted for each bidder. A feasible solution to WDP is called a *feasible allocation*. When bids are true values, i.e., $b = v$, an optimal solution is called *efficient*. An optimal solution is said to be *efficient with respect to bids* b where necessary to avoid confusion.

For a selected optimal solution x^*, bidder i's *allocated bundle* will be denoted a^{i*}, equal to the unique column of $\boldsymbol{A}$ with $x^*_k = 1$ and $k \in \mathcal{K}_i$ or the zero vector/empty bundle $\mathbf{0}$. (Note that some literature uses an x or X variable to denote an *allocation* directly; here we use x as a vector of binary *selection* variables to choose each bidder's allocation from among fixed coefficient columns a^k, emphasizing both an integer programming (IP) convention and that in practice, with XOR, bidders submit bundles which are then treated as parameters.) The set of bidders receiving a nonempty bundle at x^* will be denoted $\mathcal{I}^*$, with winning bids denoted $\mathcal{K}^*$.

The linear relaxation of WDP replaces (4) with $\mathbf{0} \leq x \leq \mathbf{1}$, and will be called WDP-LP. Because a solution to WDP-LP is potentially fractional, we write $w^f(\mathcal{I}, \mathcal{J}, \mathcal{K})$ for its objective function. Similarly, x^f is a solution to WDP-LP, used when it is different from x^*.

3.2. Payments, Prices, and Equilibrium

In a sealed-bid CA, the auctioneer assigns an allocated bundle a^{i*} and a monetary *payment* amount ρ_i to each bidder i. We assume *quasilinear utility*: a bidder's payoff or utility is given by $u_i(a^{i*}, \rho_i) = v_i(a^{i*}) - \rho_i$. Such an auction mechanism is called *individually rational (IR)* whenever $\rho_i \leq v_i(a^{i*})$ for all i, including $\rho_i = 0$ when $a^{i*} = \mathbf{0}$. Similarly, the auction is *IR with respect to bids* when $\rho_i \leq b_i(a^{i*})$ for all i.

Of particular importance are *linear anonymous* prices, where payments can be computed as $\rho_i = pa^{i*}$ for some row vector p with linear anonymous prices $p_j \in \mathbb{R}_{\geq 0}$ for each item j. When a linear

anonymous price vector p is announced, each bidder can compute her *demand correspondence*, defined as:

$$D_i(p) = \arg\max_{a \le c} v_i(a) - pa$$

This is the bidder's set of preferred bundles at the announced prices, with emphasis that this set could contain multiple bundles, including perhaps the empty bundle.

A solution to (WDP) and price vector p are said to be *envy-free* whenever each bidder prefers her allocated bundle at the current prices, i.e., when $a^{i*} \in D_i(p)$. An efficient, envy-free allocation-price pair are defined to be a *Walrasian equilibrium* (WE) if also every item j with excess supply has $p_j = 0$. This latter condition can be represented as a complementary slackness condition in vector form: $p(c - \boldsymbol{A}x^*) = 0$.

3.3. Computing Dual Prices

In this paper, we explore finding AWEs, which will act like standard WEs when artificially added items are treated as real items. Accordingly, if we assume $\boldsymbol{A}$ contains a final set of added cuts, WDP-LP will solve to integer optimality, and dual prices form a WE per the main result of Bikhchandani and Mamer (1997). The dual of WDP-LP is:

$$\begin{aligned} \min \quad & pc + s\mathbf{1} && \text{(WDP-dual)} \\ \text{s.t.} \quad & p\boldsymbol{A} + s\boldsymbol{B} \ge b && (5) \\ & p \in \mathbb{R}^J_{\ge 0}\ , \quad s \in \mathbb{R}^I_{\ge 0} && (6) \end{aligned}$$

Each constraint (5) is of the form $\sum_{j \in \mathcal{J}} a^k_j p_j + s_{i(k)} \ge b_k$, thus enforcing envy-freeness for non-winning bids. When WDP-LP is integral at optimality, $s_{i(k)} = b_k - \sum_{j \in \mathcal{J}} a^k_j p_j$ for each winning bid. Thus, dual variables s_i are the *surplus* amount for each bidder i, the *apparent* utility, treating the bid as true value and subtracting all relevant prices in the bundle, with $s_i = u_i$ when $b = v$.

Solutions to WDP-dual are usually not unique, and often in the worst case for bidders, high prices giving $s_i = 0$ for all i are feasible. Thus, we will optimize over the set of solutions to instead ensure *bidder-optimal payments*, as discussed in Ausubel and Milgrom (2002), with the motivation of providing better incentives to bidders. While there may be a number of ways to select bidder-optimal payments, we provide one definitive selection process as:

$$\begin{aligned} \min \quad & p\boldsymbol{C}p^T && \text{(WDP-quad-dual)} \\ \text{s.t.} \quad & (5) - (6) && \\ & pc + s\mathbf{1} = bx^* && (7) \end{aligned}$$

where (7) locks in the known optimal objective value to WDP-dual from a given solution x^* to WDP, and where $\boldsymbol{C}$ is a $J \times J$ matrix with diagonal $\boldsymbol{C}_{j,j} = c_j$ and all other entries zero. This

implements a quadratic minimization of squared prices (for c_j copies of each item) over all solutions to WDP-dual.

4. Pricing Valid Cuts as Artificial Items

There are many methods for adding cuts to an IP whose LP relaxation is non-integral. Indeed there is a vast literature on the subject, generally dealing with the development of practical algorithms for solving IPs. We utilize a few of the tools from that literature stream (e.g., Chvátal-Gomory cuts and lifting), but here, despite the NP-hardness of solving WDP in general (see, for example, Lehmann et al. 2006), we assume that we are working on problems of a practical size, ones for which we can comfortably solve the NP-hard WDP with available software, such as CPLEX or Gurobi.

4.1. Motivation and examples

Before presenting the details and properties of PME pricing, we show a few small examples to build intuition. We begin with a demonstration of the non-existence of WE prices for CAs.

Example 1. $b_1(AB) = b_2(AC) = b_3(BC) = 10; \quad \boldsymbol{b_4(ABC) = 12}$

Here, we have $I = K = 4$ single-minded bidders/bids, each bidding on exactly one bundle from a set of $J = 3$ unique items, $\{A, B, C\}$. Bid 4 is the only winning bid in the unique optimal solution to WDP. (We will indicate winning bids in these examples with ***bold*** font.) No IR prices constitute a WE, because envy-freeness demands:

$$p_A + p_B \geq 10, \quad p_A + p_C \geq 10, \quad \text{and} \quad p_B + p_C \geq 10$$

which combine to yield $p_A + p_B + p_C \geq 15$, a violation of IR with $15 > 12$. It is easy to show that the associated WDP has a *valid cut*, an inequality that maintains all integer solutions but may exclude fractional solutions to WDP-LP:

$$x_1 + x_2 + x_3 + x_4 \leq 1$$

Being a valid cut, adding this new row to $\boldsymbol{A}$ will not change the set of feasible solutions to WDP. Such an addition is equivalent to considering a similar auction given by $b_1(ABD) = b_2(ACD) = b_3(BCD) = 10$; $b_4(ABCD) = 12$, i.e., by adding a new item D to every bundle. In this equivalent auction, WE prices are now given by $p = [0, 0, 0, 10]$, which can be found using WDP-quad-dual.

Intuitively, the cut says, only one of these four bids can win. Thus, if we were to make a license D and require that a bidder must buy D to be eligible to win at least two items from $\{A, B, C\}$, the outcome would be equivalent to a second-price auction for D and awarding $\{A, B, C\}$ each at

zero (additional) cost. Though the auctioneer may have anticipated this and structured the items for sale to include such a license D, the point here is that she didn't have to; the auction format we propose finds and exploits the relevant structure to produce a WE, learning about the market and how to effectively price the items (including by possibly adding valid artificial items) based on the bids.

One might also notice that the new WE outcome with $p_D = 10$ is equivalent to the Vickrey-Clarke-Groves (VCG) auction outcome on the original bids. While our cut procedure does provide a price linearization of the VCG outcome on some problems like this one, this is not always the case, as we will see below with Example 5 and Corollary 3.

Example 2. $\boldsymbol{b_1(9A) = 85}$; $b_2(4A) = 40$; $b_3(4A) = 40$; $b_4(4A) = 35$;

Here, with a supply of $c_A = 9$ units of a single type, and again single-minded bidders, $b_2(4A) = 40$ dictates that $4p_A \geq 40$ and therefore $p_A \geq 10$. But this cannot be satisfied with IR constraint $9p_A \leq 85$, so no (natural) WE is possible.

By multiplying the lone row of $\boldsymbol{A}$ by 0.25 and rounding down non-integer coefficients (an instance of the Chvátal-Gomory process), we find the cut $2x_1 + x_2 + x_3 + x_4 \leq 2$, which can be priced $p = [0, 40]$, resulting in AWE. Here, the added cut can be thought of as a fictional item packaging 4 items together, because the losing bids indicate that this is the minimum size of a valuable package. The winning bid on 9 items therefore must buy two copies of this packaged quartet of items (reflected by the cut coefficient of 2 for x_1) and receive one original ("unpackaged") item free of charge, altogether paying $\rho_1 = 80$.

Example 3. $b_1(AB) = 18$; $b_2(BC) = 17$; $\boldsymbol{b_3(CD) = 18}$;
$b_4(DE) = 16$; $\boldsymbol{b_5(AE) = 19}$; $b_6(ABCDE) = 34$

Again with single-minded bidders and unique items, here only two of the first five bundles could ever be accepted (considering only feasibility of item bundles), or, the sixth bundle could be accepted and knock out any other two bids. This statement is not difficult to read intuitively from the generated cut:

$$x_1 + x_2 + x_3 + x_4 + x_5 + 2x_6 \leq 2$$

Though no WE exists before the cut, with the cut we have a WE with $p = [2, 0, 1, 0, 0, 16]$, resulting in $\rho_3 = 17$ and $\rho_5 = 18$.

Of interest, unlike the previous examples, this equilibrium uses non-zero prices on some original items in addition to pricing the generated artificial item. Items D and E are in the lesser valued losing bid 4, which sets the value of the added cut item at 16. This is the minimum price to take

two items from $\{A, B, C, D, E\}$. With $p_B = 0$ in a WE because B does not sell, setting $p_A = 2$ and $p_C = 1$ covers the entire "upsell" above 16 to satisfy envy-freeness on $b_1(AB)$ and $b_2(BC)$. Bid 6 would face a price of 35 (the price of A, C, and 2 copies of the added item, per the added cut) which she does not envy. This cut demonstrates the "odd-hole" property discussed in the valid-cut literature.

4.2. Price-Match Equilibria for CAs

While it is possible to add just enough cuts to guarantee WE, we propose going further, with deeper cuts yielding what we call a *Price-Match Equilibrium* (PME) which we now motivate.

Adomavicius and Gupta (2005) define a *winning level* of a bundle a at any instant in a continuous efficient CA as the amount needed for a new bid on a to just become winning, given the bids submitted up until that point in time. They show that it can be computed in our notation (before the new bid has been added to $\mathcal{K}$) as $w(\mathcal{I}, \mathcal{J}, \mathcal{K}) - w(\mathcal{I}, \mathcal{J} \setminus a, \mathcal{K})$, where $\mathcal{J} \setminus a$ represents the same item classes but with supply vector reduced to $c - a$.

Each payment in Examples 1, 2, and 3 is a winning level under this definition. For instance, in Example 1, remove the winning bid $k = 4$ and find the level needed to just barely becoming winning (or to become strictly winning if just ϵ higher). compute $w(\mathcal{I} \setminus 4, \mathcal{J}, \mathcal{K} \setminus 4) = 10$ and $w(\mathcal{I} \setminus 4, \mathcal{J} \setminus ABC, \mathcal{K} \setminus 4) = 0$, and thus $\rho_4 = 10 - 0$, consistent with the AWE computed for this example. Despite the prices having the winning-level property in these first examples, this is not true in general.

Proposition 1. *Minimal Walrasian Equilibrium prices may be strictly greater than winning-level payments.*

The proof is by example. We alter Example 1 by raising the winning bid amount from 12 to 16, relaxing $b_4(ABC)$'s IR constraint:

Example 1.1: $b_1(AB) = b_2(AC) = b_3(BC) = 10$; $\quad \boldsymbol{b_4(ABC) = 16}$

Now with a higher winning bid, WDP and WDP-LP no longer have an integrality gap. WDP-quad-dual prices $p = [5, 5, 5]$ give a WE with $\rho_4 = 15$ for Example 1.1 (compared to $\rho_4 = 10$ in Example 1, after adding a cut). Note that p is minimal in the available payment space, i.e, with no cut and pricing only three items. This provides strong motivation to add a cut, even when there is no integrality gap; the seller has no feasible way to generate revenue of more than 10 without bidder 4, so why should bidder 4 be forced to pay 50% more? Applying the same cut as in Example 1 to Example 1.1, results in the same $\rho_4 = 10$, which again is the VCG outcome.

Indeed, in Example 1 , without the cut, the first three bids would be "half accepted" by solution $x^f = [0.5, 0.5, 0.5, 0]^T$ to WDP-LP. In doing so, it's as if the three bids are allowed to represent 15 units of value, when really they only generate 10 units of value integrally. In Example 1.1 (if a cut is not applied), the raised bid "hides" this non-convex over-influence of these three bids, allowing them to drive prices up beyond their ability to match those prices in any integral solution.

To further refine the winning-level notion for the current context (i.e., a multi-unit, sealed-bid CA with linear prices and multi-minded bidders), consider the following definition:

Definition 1 (Price Match Equilibrium). *A WE allocation x^* and price vector p constitute a* Price-Match Equilibrium *(PME) if and only if $\forall i$, there exists an allocation $x^{-i} \in \{0,1\}^K$ with $x_k^{-i} = 0$ for $k \in \mathcal{K}_i$ and for which the seller could generate the same total payments (i.e., for which $p\boldsymbol{A}x^{-i} = p\boldsymbol{A}x^*$) subject to IR with respect to bids.*

Next, let $K(p)$ represent the original bids but with amounts b changed to b^p where $b_k^p = pa^k$ if $pa^k \leq b_k$, $b_k^p = 0$ otherwise, i.e., $K(p)$ replaces each bid amount with zero or the total linear price for the associated bundle a^k according to p whenever that price would yield non-negative surplus. We then have the following:

Proposition 2. *WE allocation-price pair (x^*, p) is a PME if and only if:*

$$\rho_i = pa^{i*} = w(\mathcal{I} \setminus i, \mathcal{J}, \mathcal{K}(p)) - w(\mathcal{I} \setminus i, \mathcal{J} \setminus a^{i*}, \mathcal{K}(p)) \qquad \forall i \in \mathcal{I}$$

According to this result, in a PME each winning bidder's payment ρ_i is (simultaneously) in the functional form of a winning level; each i sees prices that the other bidders are willing to pay, both with and without her own consumption of a^{i*}, and thus that her own payment if lowered would not be winning. For any CA with IR, the condition is trivially satisfied for all non-winning i paying $\rho_i = 0$. The WE outcomes of Examples 1, 2, and 3 can easily be shown to be PMEs, while Example 1.1 with $p = [5, 5, 5]$ is a WE but not a PME.

Directly from the definition of PME, the second-price auction for a single item has a PME outcome. Similarly, the Generalized Second-Price Auction, in which the winner of the nth slot pays the amount of the $n+1$st bid (famous for its use in Google's ad-word auctions), prescribes a PME outcome: if any bidder were to leave, each bidder in a worse-ranked slot could be moved up one slot without violating IR. Considering the more general CA environment, per de Vries et al. (2007), when coalitional submodularity holds, VCG is implementable as a NWE. Together with Corollary 3 below, we have:

Proposition 3. *Coalitional submodularity is sufficient to guarantee min-NWE = PME = MRC = VCG.*

where min-NWE is the payment-minimal NWE, which is unique in that setting.

In a WE, the seller prefers to sell the efficient allocation, here represented as: $p\boldsymbol{A}x^* = w(\mathcal{I},\mathcal{J},\mathcal{K}(p))$. From WDP's definition, removing a bidder yields $\sum_{\bar{\imath}\in\mathcal{I}\setminus i}\rho_{\bar{\imath}} = w(\mathcal{I}\setminus i,\mathcal{J}\setminus a^{i*},\mathcal{K}(p))$, which added to the condition of Proposition 2 yields an alternate characterization of PME:

Proposition 4. *WE allocation-price pair (x^*,p) is a PME if and only if:*

$$p\boldsymbol{A}x^* = w(\mathcal{I}\setminus i,\mathcal{J},\mathcal{K}(p)) \qquad \forall i\in\mathcal{I}$$

5. Comparison to Min-Revenue Core-Selecting and VCG Mechanisms

Using our notation, truth-inducing VCG payments are computed as:

$$\rho_i^{VCG} = w(\mathcal{I}\setminus i,\mathcal{J},\mathcal{K}) - w(\mathcal{I}\setminus i,\mathcal{J}\setminus a^{i*},\mathcal{K}) \tag{8}$$

Minimum-Revenue Core-Selecting (MRC) auctions minimize $\sum_{i\in\mathcal{I}}\rho_i$ such that:

$$\sum_{i\in\mathcal{C}}(b_{i*}-\rho_i) + \sum_{i\in\mathcal{I}}\rho_i \geq w(\mathcal{C},\mathcal{J},\mathcal{K}^{\mathcal{C}}) \quad \forall\mathcal{C}\subseteq\mathcal{I} \tag{9}$$

where $\mathcal{K}^{\mathcal{C}}$ is shorthand for $\cup_{i\in\mathcal{C}}\mathcal{K}_i$.

Proposition 5. *Any WE pair (x^*,p) is core-selecting.*

The proof sketch is that core-selecting constraints as in Equation 9 can be derived by simply adding the relevant envy-freeness constraints (5) indexed by winning bids in the $w(\mathcal{C},\mathcal{J},\mathcal{K}^{\mathcal{C}})$ outcome. The converse does not directly hold: not all core outcomes can be reached via the currently available items prices. (This will be addressed later by Theorem 1, however.) To demonstrate this important difference between MAP prices and the VCG and MRC benchmarks from the literature, consider the following numerical example, adapted from Beck (2013):

Example 4. $\boldsymbol{b_1(AC)=21}$; $b_1(A)=10$; $b_1(C)=11$;
$b_2(BC)=20$; $\boldsymbol{b_2(B)=10}$; $b_2(C)=10$;
$b_3(AB)=10$

Beck (2013) demonstrates a weakness of MRC, referred to as the "unrelated goods problem," motivated by the observation that bidders 1 and 2 have preferences for C that simply add to their values for either A or B, respectively. An auctioneer considering holding separate MRC auctions for $\{A,B\}$ and $\{C\}$, or selling $\{A,B,C\}$ together would determine:

$\rho_1^{MRC}=5$ for A	$\rho_2^{MRC}=5$ for B	when selling $\{A,B\}$
$\rho_1^{MRC}=10$ for C		when selling $\{C\}$
$\rho_1^{MRC}=10$ for AC	$\rho_2^{MRC}=0$ for B	when selling $\{A,B,C\}$

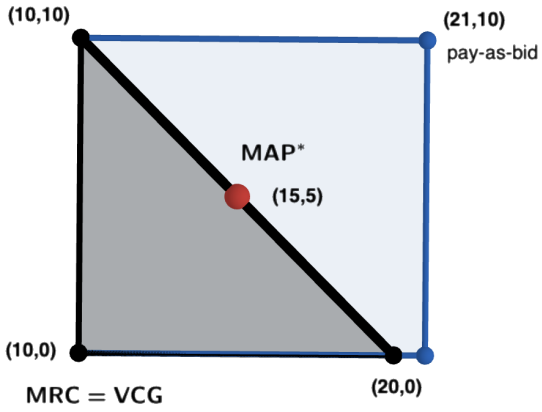


Figure 2 **The core of Example 4 in the ρ_1,ρ_2 plane. The lighter region consists of NWE payment outcomes while the AWE = PME outcomes form the gray shaded triangle. The thicker line segment intersection of these two regions is the image of MAP prices on the core. WDP-quad-dual selects MAP* (red point) from among the PME points. The pay-as-bid and MRC = VCG points are also labeled.**

and observe that total revenue goes down drastically (from 20 to 10) when combining the $\{A, B\}$ auction with that of an "unrelated good" $\{C\}$ into a single auction. Beck (2013) notes that this shows a similar "low-revenue" problem as VCG, observing that total payments collected by VCG are 10 whether these two auctions are conducted separately or together. With this example as motivation, we observe the unrelated goods problem vanish for MAP prices as proposed here.

Indeed, for Example 4, we have WDP-quad-dual prices $p = [5, 5, 10]$ forming a PME without the addition of any artificial cut items. Further, we have the same prices supporting two auctions if $\{A, B\}$ and $\{C\}$ are conducted separately, with total revenue of 20 for the seller in either case. One might also observe that any p with $p_A + p_B = 10$ and $p_c = 10$ are sufficient for PME here, or that any $\rho_1^{MRC} + \rho_2^{MRC} = 10$ satisfy MRC when selling just $\{A, B\}$; this motivates our use of quadratic rules to specify a unique pricing, to select exactly MAP*, as shown in Figure 2.

Intuitively, because MRC only demands personalized payments, in Example 4 bidders 2 and 3 cannot "tell the difference" between payments for A and payments for C, and thus a form of "double counting" occurs. Bidder 2 sees bidder 1's MRC payment of 10 and would not (in a hypothetical sense) attempt to block the outcome and pay the seller more, as if bidder 1's payment was for C. But at the same time, bidder 3 sees bidder 1's payment and also cannot block, as if the payment of 10 was for A. By accounting for the prices of items separately, MAP prices avoid this double counting, and thus remain greater than VCG = MRC for this example.

6. Generating Cuts for WE and PME

We first consider generating a single valid cut. Let α_0 be the right-hand side of a new cut, with cut coefficients α_k for each $k \in \mathcal{K}$, the latter collectively forming the vector α, together forming a cut:

$$\alpha x \leq \alpha_0$$

which must hold for any WDP-feasible allocation $x \in \{0,1\}^K$. These α terms will be variables during the cut generation process and then later treated as constants in the final price determination, functioning exactly as all other items in $\mathcal{J}$, and concatenated as a new row of the WDP matrix $\boldsymbol{A}$. This new cut will thus be treated as an *artificial item* in the auction. Eventually, we prefer a collection of artificial items with each $(\alpha, \alpha_0) \in \mathbb{Z}^{K+1}_{\geq 0}$ for interpretability, but for intermediate steps in our reasoning and algorithms we will also consider "unscaled" or continuous valid cuts, which we denote with a bar-notation: $(\bar{\alpha}, \bar{\alpha}_0) \in \mathbb{R}^{K+1}_{\geq 0}$.

Any newly added artificial item must be *fully demanded* at our fixed WDP solution x^*, i.e., $\alpha x^* = \alpha_0$. That is every copy of an artificial item must be purchased by the winning bids. If not, WE fails by definition, and pricing an excess-supply item may violate the core property, even when envy-freeness constraints (5) are satisfied. Because the preponderance of previous work on cut generation was used to *solve* IPs like WDP, this assumption that we *know* x^* and constrain that $\alpha x^* = \alpha_0$ when looking for cuts is peculiar to the current line of research, to our knowledge.

6.1. Cuts for WE

Though it might be computationally expensive in practice, one can in theory enumerate all set-wise maximal feasible solutions to WDP, those for which no other bid could be feasibly added. Let all such maximal feasible solutions be indexed by $\ell \in \mathcal{L}$ with elements x^ℓ. If we construct the matrix of directions to optimality, $\boldsymbol{D} = [x^* - x^1, x^* - x^2, ..., x^* - x^{|\mathcal{L}|}]$, then one direct approach to find a cut (α, α_0) is by solving:

$$\begin{aligned} \max \quad & \alpha x^f && \text{(cut-frac)} \\ & \alpha \boldsymbol{D} \geq \boldsymbol{0} && (10) \\ & \alpha x^* = \alpha_0 && (11) \end{aligned}$$

where x^f is a fixed fractional solution to WDP-LP, treated here as a constant, as is $\boldsymbol{D}$. (If WDP-LP admits an integral solution, we are at a WE, and so we would terminate without the need for a new cut.) While constraints (10) ensure that (α, α_0) is a valid cut for WDP, any positive scalar multiple is also valid, resulting in an unbounded objective. To avoid this unboundedness, one can begin with $\alpha_0 = 1$ and later scale (α, α_0) to find an integer cut, as described in Algorithm 1.

Algorithm 1 Heuristic to Find AWE Prices

Solve WDP and WDP-LP

while $w^f(\mathcal{I},\mathcal{J},\mathcal{K}) > w(\mathcal{I},\mathcal{J},\mathcal{K})$ **do**

Solve (cut-frac) with $\bar{\alpha} \in \mathbb{R}^K_{\geq 0}$ and fixed $\bar{\alpha}_0 = 1$

Find the smallest $\gamma \geq 1$ such that $\forall k \in \mathcal{K}$, $\gamma\bar{\alpha}_k \in \mathbb{Z}$

Update $(\alpha, \alpha_0) \leftarrow (\gamma\bar{\alpha}, \gamma\bar{\alpha}_0)$

Add cut (α, α_0) to all formulations

Re-solve WDP-LP, updating $w^f(\mathcal{I},\mathcal{J},\mathcal{K})$ and x^f

Solve WDP-quad-dual (including all cuts as new artificial items)

Proposition 6. *Algorithm 1 converges to AWE prices for any CA instance.*

As a sketch of proof, note that a maximally violated cut maximizes $\alpha x^f - \alpha_0$, so that (cut-frac) finds a maximally violated cut at each iteration, given a fixed α_0. Progress therefore must be made at each cut iteration, lowering the WDP-LP objective value, leading to eventual convergence. The scaling integer γ is the least common multiple of denominators of reduced rational α_k values, updating them to integer values for a more interpretable artificial item, motivated by our previous examples.

Solved as an LP (i.e., with no integrality constraints within the while loop), Algorithm 1's cut generation appears easy; the difficulty is in generating the large matrix $\boldsymbol{D}$, requiring the enumeration of solutions to WDP, which is NP-Hard in general. For small examples, this is not problematic, but as the auction becomes larger, columns of $\boldsymbol{D}$ will need to be dynamically generated to manage the size. (Details on the dynamic generation of $\boldsymbol{D}$ are described in Appendix A.) Alternatively, since it is clear from the related literature that there are often many ways to find cuts to a non-integral IP formulation, we provide Algorithm 1 as just one direct cutting algorithm for completeness, leaving the development of faster or more specialized techniques for future research.

6.2. Cuts for any Core Point and PME

Example 1.1 demonstrated that a WE may not be a PME, motivating us to consider additional cuts to WDP to get to PME. Algorithm 1 can take us no further in generating new cuts once we have an AWE (or if we start at NWE), as its progress was driven by closing an IP gap that is now closed. For another, perhaps more interesting illustration, consider the single-minded bidders in the Example 5 that follows. Here there are $c_A = 17$ copies of one item, A. (This problem is adapted from an example on lifting to get cuts to the knapsack problem from Bertsimas and Weismantel 2005.)

Example 5. $\boldsymbol{b_1(5A) = 100}$; $\boldsymbol{b_2(5A) = 100}$; $\boldsymbol{b_3(5A) = 100}$; $b_4(3A) = 25$; $b_5(8A) = 90$;

Here, there is clearly a fractional solution, because only 15 of the 17 items sell, and a losing bid will thus activate fractionally in WDP-LP. There is also an easy-to-infer cut, $x_1+x_2+x_3+x_4+x_5 \leq 3$, i.e., "There can be at most 3 winning bids." This cut gives an AWE with $p=[0,90]$, but these fail to be a PME. (If one winner leaves, there is no way to get back to total payments of 270.)

A second cut can be found, $x_1+x_2+x_3+2x_5 \leq 3$, reflecting, "Bid 5 knocks out two of the three winning bids." Pricing this inferred fact as a second cut gets us PME prices $p=[0,25,32.5]$. Intuitively, bid 4 sets a minimum value of 25 to be among the winners, and then an additional 65 is required from any combination of two winners so that bid 5 will remain envy-free. One may notice that these correspond to symmetric MRC payments $\rho_i^{MRC}=57.5$, with tight core constraints that say in words, "Any two winners' payments match the combined bids 4 and 5 which total 115." One may also note that $\rho_i^{VCG}=25$ is not in the core for this example.

We observe in Example 5 that one may generate an AWE that is not a PME, but then cut further to get a PME. Further, in Example 1.1 we observe that a natural WE can be cut to a PME. From these observations it is natural to ask, what happens more generally if we continue to make cuts past the point of establishing WE prices? To characterize the closure of cut generation for WDP, we next show that any core point can be reached as a feasible AWE by adding at most one cut.

Theorem 1. *For any core-selection $(\bar{\rho}, x^*)$ there exists an AWE with total payments equal to $\bar{\rho}_i$ for each winning bidder i.*

Proof. For a proof by contradiction, suppose that the following LP has no solution:

$$\begin{aligned}
\min \quad & \alpha \mathbf{0} & \\
\text{s.t.} \quad & pa^k + \alpha_k = \bar{\rho}_{i(k)} & \forall k \in \mathcal{K}^* \\
& pa^k + \alpha_k \geq b^k - b^{i(k)*} + \bar{\rho}_{i(k)} & \forall k \in \mathcal{K} \setminus \mathcal{K}^* \\
& \alpha \boldsymbol{D} \geq \mathbf{0} & \\
& p \in \mathbb{R}^J_{\geq 0}\ , \quad \alpha \in \mathbb{R}^K_{\geq 0} &
\end{aligned}$$

By construction, the existence of a solution would be an adjusted set of prices p that equal core payments $\bar{\rho}_i$ for each i while maintaining WE, and are made by a valid cut $(\alpha, \alpha x^*)$, according to the three sets of linear constraints, respectively. When this formulation is infeasible, its dual must be unbounded, noting that it is feasible when all variables are zero. This dual can be written:

$$\max \quad \sum_{k \in \mathcal{K}^*} \bar{\rho}_{i(k)} y_k + \sum_{k \in \mathcal{K} \setminus \mathcal{K}^*} (b^k - b^{i(k)*} + \bar{\rho}_{i(k)}) z_k$$

$$\text{s.t.} \quad \sum_{k\in\mathcal{K}^*} a_j^k y_k + \sum_{k\in\mathcal{K}\setminus\mathcal{K}^*} a_j^k z_k \le 0 \qquad \forall j\in\mathcal{J} \tag{12}$$

$$\sum_{\ell\in\mathcal{L}} \beta_\ell (1-x_k^\ell) \le -y_k \qquad \forall k\in\mathcal{K}^* \tag{13}$$

$$z_k \le \sum_{\ell\in\mathcal{L}} \beta_\ell x_k^\ell \qquad \forall k\in\mathcal{K}\setminus\mathcal{K}^* \tag{14}$$

$$y\in\mathbb{R}_{\le 0}^{|\mathcal{K}^*|}\ , \quad z\in\mathbb{R}_{\ge 0}^{|\mathcal{K}\setminus\mathcal{K}^*|}\ , \quad \beta\in\mathbb{R}_{\ge 0}^{|\mathcal{L}|}$$

with x_k^ℓ terms constant in this formulation from the $\boldsymbol{D}$ matrix. Note that y variables are non-positive as a consequence of (13). Continuous y variables select negative amounts of winning bids, while z variables select positive amounts of non-winning bids. Any increase of a z variable must be offset by at least as much change among y variables in the opposite direction, on an item-by-item basis per constraints (12). Associated with each $\ell\in\mathcal{L}$ is a maximal feasible solution x^ℓ; let $\mathcal{K}^\ell=\{k\in\mathcal{K}|x_k^\ell=1\}$. Because $\bar{\rho}$ is core-selecting:

$$\sum_{k\in\mathcal{K}^*\setminus\mathcal{K}^\ell} \bar{\rho}_{i(k)} \ge \sum_{k\in\mathcal{K}^\ell\setminus\mathcal{K}^*} (b^k - b^{i(k)*} + \bar{\rho}_{i(k)})$$

Otherwise, ℓ would provide a blocking coalition at $\bar{\rho}$. Multiplying each of these by β_ℓ and summing over $\mathcal{L}$ produces:

$$\sum_{k\in\mathcal{K}^*} \bar{\rho}_{i(k)} (\sum_{\ell\in\mathcal{L}} \beta_\ell (1-x_k^\ell)) \ge \sum_{k\in\mathcal{K}\setminus\mathcal{K}^*} (b^k - b^{i(k)*} + \bar{\rho}_{i(k)}) (\sum_{\ell\in\mathcal{L}} \beta_\ell x_k^\ell)$$

Together with (13) and (14), we have:

$$-\sum_{k\in\mathcal{K}^*} \bar{\rho}_{i(k)} y_k \ge \sum_{k\in\mathcal{K}\setminus\mathcal{K}^*} (b^k - b^{i(k)*} + \bar{\rho}_{i(k)}) z_k$$

But this implies that the maximization objective is always non-positive, contradicting its supposed unboundedness. □

Theorem 1 ensures that all (personalized) core payments can be reached as a linearization via *some* valid cut. In Example 4, NWE prices were found to be PME prices, in the core by Proposition 5, but distinct from VCG = MRC payments. To arrive at the latter, an intuitive cut $x_1 + x_4 + x_7 \le 1$ could be added (giving a k index to each bid in the order listed), linearizing the VCG = MRC outcome with $p = [0, 0, 0, 10]$, pricing the added cut as a "price of 10 to have a two-item bundle." This cut was not needed to establish a PME, motivating us to only add cuts as strong as needed for a PME, and not further to MRC, lest we invoke the unrelated goods problem. MRC payments are useful here, however, as minimal PMEs, helping establish that PMEs always exist and are obtainable through valid cuts:

Theorem 2. *If WE prices p are not a PME, there is a core point with lower total payments. Specifically, $p\boldsymbol{A}x^{-\bar{\imath}} < p\boldsymbol{A}x^*$ implies $(\bar{\rho}, x^*)$ is in the core, where $\bar{\rho}_i = pa^{i*}$ for $i \neq \bar{\imath}$ and $\bar{\rho}_{\bar{\imath}} = pa^{\bar{\imath}*} - \epsilon$ for some small $\epsilon > 0$.*

Proof. Suppose that for some $\bar{\imath}$ the specified $\bar{\rho}$ is not in the core for any ϵ. Let ρ and s be payments and surplus under p. With ρ in the core by Proposition 5 but $\bar{\rho}$ not in the core after an ϵ reduction of bidder $\bar{\imath}$'s payment, it must be the case that there is some tight core constraint at ρ, indexed by $\mathcal{C}$ with $\bar{\imath} \notin \mathcal{C}$:

$$\sum_{i \in \mathcal{I}^* \setminus \mathcal{C}} \rho_i = w(\mathcal{C}, \mathcal{J}, \mathcal{K}^{\mathcal{C}}) - \sum_{i \in \mathcal{C}} b_{i*}$$

Let each winning bid amount in $w(\mathcal{C}, \mathcal{J}, \mathcal{K}^{\mathcal{C}})$ be given by $b_i^{\mathcal{C}}$ for bundle $a^{\mathcal{C},i}$ for each $i \in \mathcal{C}$, with selection vector $x^{\mathcal{C}}$. Substitute these bid amounts $b_i^{\mathcal{C}}$ into the previous equation for $w(\mathcal{C}, \mathcal{J}, \mathcal{K}^{\mathcal{C}})$ and then substitute $b_{i*} = s_i + \rho_i$, moving ρ_i terms to the LHS to get:

$$\sum_{i \in \mathcal{I}^*} \rho_i = \sum_{i \in \mathcal{C}} (b_i^{\mathcal{C}} - s_i)$$

whose LHS is $p\boldsymbol{A}x^* =$ by definition. Next, because p supports a WE, we have $pa^{\mathcal{C},i} \geq b_i^{\mathcal{C}} - s_i$ for each $i \in \mathcal{C}$, summing to yield:

$$p\boldsymbol{A}x^* \leq p\boldsymbol{A}x^{\mathcal{C}}$$

If we have strict $pa^{\mathcal{C},i} > b_i^{\mathcal{C}} - s_i$ for any $i \in \mathcal{C}$, then this last statement is also strict, contradicting (x^*, p) as a WE. (An efficient x^* must be in the seller's preferred supply set at WE, therefore an alternative $x^{\mathcal{C}}$ with $p\boldsymbol{A}x^* < p\boldsymbol{A}x^{\mathcal{C}}$ is a contradiction.) If instead none of those envy-freeness constraints are strict, we have each $s_i = b_i^{\mathcal{C}} - pa^{\mathcal{C},i} \geq 0$, and thus $x^{\mathcal{C}}$ is an IR allocation among $\mathcal{I} \setminus \bar{\imath}$ that generates as much price revenue as x^*, contradicting the assumption that prices p are not a PME, as specifically witnessed by $\bar{\imath}$. □

Corollary 1. *If WE prices are MRC, they are also PME prices.*

If not, Theorem 2 would imply a direction to reduce prices within the core, contradicting the minimum-revenue assumption. Altogether, because MRC prices always exist and Theorem 1 says they can always be supported as AWE, we have:

Corollary 2. *For any CA, PME prices can be generated via cuts.*

Note that this result is limited to CAs, not extending beyond the assumptions of this paper to combinatorial exchanges (with endogenous supply) as the core and hence MRC is not guaranteed to exist in such a general setting. As an additional observation, using the Ausubel and Milgrom (2002) characterization of VCG being the unique MRC point when in it is in the core, we find:

Corollary 3. *VCG payments can be linearized to PME prices if and only if VCG = MRC.*

Also interestingly, despite VCG's uniqueness in MRC whenever in the core, we have the following proposition, proven by the example that follows it.

Proposition 7. *The linearization of VCG in Corollary 3 is not necessarily unique.*

Example 6. $\boldsymbol{b_1(4A) = 70}$; $\boldsymbol{b_2(4A) = 60}$ $b_3(2A) = 35$; $b_4(2A) = 20$; *with* $c_A = 9$ *copies of a single item* A.

Again, excess supply forbids a natural WE for Example 6, but PME prices can be obtained by either the left two cuts or the right two cuts:

$$2x_1 + 2x_2 + x_3 + x_4 \leq 4 \qquad x_1 + x_2 + x_3 \leq 2$$
$$x_1 + x_2 + x_3 \leq 2 \qquad x_1 + x_2 + x_4 \leq 2$$

This induces $p = [0, 20, 15]$ for left pair or $p = [0, 35, 20]$ for the right pair, both resulting in linearizations of $\rho^{VCG} = [55, 55]$. The left pair can be seen to price pairs of A plus "only two of bids 1, 2, 3 can win," while the right pair says more symmetrically, "only two of bids 1, 2, k can win" for $k = 3, 4$. Notice that the top left cut is the sum of the right two. In our elaboration of MAP prices below, we will find *minimal* cuts that maintain balance among price-setting scenarios, and so we will favor the right two. Final cuts in our procedure belong to the Hilbert basis of the cone of feasible cuts, and so the top left cut could not be selected as an explanation of final payments.

The following final example was adapted from Gul and Stachetti (1999) by Day (2004), and has a non-PME NWE and artificial PMEs with revenues strictly greater than VCG = MRC:

Example 7. *Using an OXS (OR-of-XOR-of Singletons) bid format for compactness: for each bidder, a bid may be accepted for an item from* $\{A, B, C, D\}$ *from either column, or both, but using at most one bid per column. For example, using the following bids,* $b_1(A) = b_1(C) = 8$ *and* $b_1(AC) = 16$, *but* $b_1(AB) = 8$ *because only one of the two bid amounts in the first (or any) column can be selected.*

Bidder 1

A	**8**	0
B	8	0
C	0	**8**
D	0	8

Bidder 2

A	6	0
B	0	**6**
C	2	0
D	0	0

Bidder 3

A	0	0
B	2	0
C	0	6
D	**6**	0

One of four possible tied efficient allocations is bolded. This example was originally devised to show that WE prices may exceed VCG payments, even when the gross substitutes property holds.

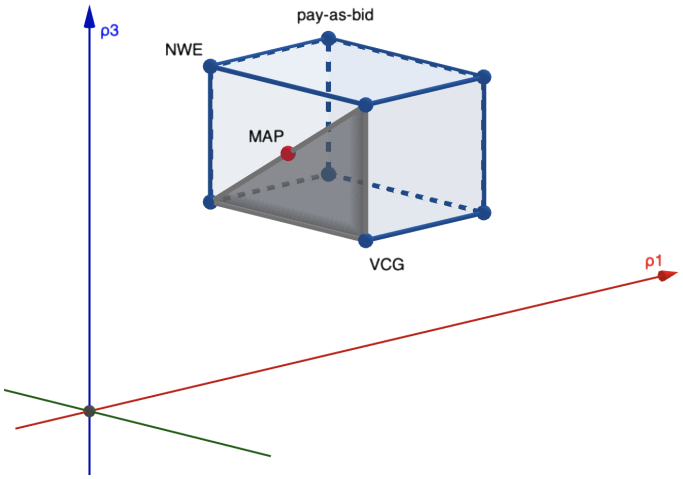


Figure 3 **The core (in payment space) for Example 7. The pay-as-bid, NWE, and VCG = MRC points are labeled. PME points form the shaded triangle on the near face, with the unique MAP point $(12, 4, 4)$ in red.**

We strengthen this to the more recently studied *strong* substitutes property, which demands gross substitutes even if items are uniquely identified. Unique NWE prices $p = [6, 6, 6, 6]$ generate 24 in revenue and are not a PME, since the seller could not sell all four items at $p_j = 6$ when either bidder 2 or 3 is removed. Payments $\rho^{VCG} = [12, 2, 2]$ generate only 16 in revenue. VCG is in the core, and thus VCG = MRC. This also shows:

Corollary 4. *The strong substitutes property, while sufficient to guarantee NWE, is not sufficient to guarantee that the NWE is a PME.*

The MAP outcome for Example 7 can be found after adding two cuts, that say in words, respectively, "Bidders 1 and 2 cannot both get a most valued bundle," and "Bidders 1 and 3 cannot both get a most valued bundle." These cuts yield $p = [4, 4, 4, 4, 2, 2]$ (the last two prices on artificial items) with total payments of 20 and bidder 1 paying for both artificial items. These payments $\rho_2 = \rho_3 = 4$ lie strictly between the other outcomes for bidders 2 and 3, but with $\rho_1 = 12$ in all three outcomes. The core of Example 7 is shown in Figure 3, including the unique MAP point in red.

6.3. Computing MAP prices

We now introduce "Minimally Artificial PME" (MAP) payments as a precise solution concept. Our computational strategy is to find a minimal vector $\bar{\alpha}$ of artificial adjustments for each bid (positive real numbers), which determines final payments. Then in a second procedure, we produce an *elaboration* of these amounts into more meaningful prices, generated as a positive linear combination of minimal cut vectors. The first optimization discovers how much artificiality *must* be introduced, and the second elaboration operation finds a *simplest* set of cuts that *explains* the artificial amounts.

The first formulation (Find-$\bar{\alpha}$) minimizes the total amount of artificial (i.e., not linear in original auction items) adjustments needed to achieve PME:

$$\begin{aligned}
\min \quad & \bar{\alpha}\mathbf{1} && && && \text{(Find-}\bar{\alpha}\text{)}\\
\text{s.t.} \quad & \bar{\alpha}\boldsymbol{D} \geq \mathbf{0} && && && (15)\\
& p\boldsymbol{A} + s\boldsymbol{B} + \bar{\alpha} \geq b && && && (16)\\
& \boldsymbol{A}^{-i}x^{-i} \leq c && \text{with } = \text{ when } p_j > 0 && \forall i \in \mathcal{I}^* && (17)\\
& \boldsymbol{B}^{-i}x^{-i} \leq \mathbf{1} && && \forall i \in \mathcal{I}^* && (18)\\
& \bar{\alpha}x^{-i} = \bar{\alpha}x^* && && \forall i \in \mathcal{I}^* && (19)\\
& pa^k + \bar{\alpha}_k \leq b_k && \text{when } x_k^{-i} = 1 && \forall i, k \in \mathcal{I}^* \times \mathcal{K} && (20)\\
& \bar{\alpha} \in \mathbb{R}^K_{\geq 0} && && && (21)\\
& s \in \mathbb{R}^I_{\geq 0} && && && (22)\\
& p \in \mathbb{R}^J_{\geq 0} && \text{with } p(c - \boldsymbol{A}x^*) = 0 && && (23)\\
& x^{-i} \in \{0,1\}^{K-K_i} && && \forall i \in \mathcal{I}^* && (24)
\end{aligned}$$

Here we present Find-$\bar{\alpha}$ compactly, with nonlinear dependencies and quadratic constraints. An equivalent (but longer) linear integer program is presented in Appendix B for completeness.

If there exists a NWE that is a PME, standard linear item prices are sufficient, and the objective will be zero, with $\bar{\alpha} = \mathbf{0}$, and thus MAP = NWE. If, on the other extreme, all original items have excess supply, linear prices cannot be used at all, and $\bar{\alpha}$ does all the work of pricing with its nonlinear terms, as occurs in Examples 5 and 6.

Constraints (15) like (10) ensure a valid cut, while (16) like (5) are WE-defining envy-freeness constraints, but here with the assistance of additional nonlinear $\bar{\alpha}$ terms. Constraints (17) and (18) ensure that when any winning bidder i is removed, x^{-i} variables select a feasible WDP solution without i, one that sells any item that receives a positive price. Set (19) ensures that (continuous) artificial item α sells completely in each x^{-i} scenario. Together with (20) that ensure IR in each case, (17)-(20) explicitly fulfill the definition of PME when each bidder is removed. Surplus variables s and thus (22) can be substituted out because x^* is known, but we leave them in this formulation to show the connection to the simple WDP-dual. The complementary slackness condition shown in (23) is easy to implement simply as $p_j = 0$ if item j has excess supply under x^*.

According to Theorem 1 and Corollary 2, formulation Find-$\bar{\alpha}$ will always have a solution, given the assumptions of this paper (a single-seller CA normalized to zero reserve prices). If these continuous MAP payment amounts $\bar{\alpha}$ are not all zero, we next perform what we will call *elaboration*

to produce MAP item *prices*. (If $\bar{\alpha} = \mathbf{0}$ no elaboration is needed.) To do so, we first introduce some additional nomenclature.

Based on the solution of formulation Find-$\bar{\alpha}$, let $\bar{\boldsymbol{D}}$ be the tight constraints from (15), i.e., those columns of $\boldsymbol{D}$ such that $\bar{\alpha}\bar{\boldsymbol{D}} = \mathbf{0}$. Call these columns $\bar{\mathcal{L}}$ with $\bar{L} = |\bar{\mathcal{L}}|$. Further, replace any row k of $\bar{\boldsymbol{D}}$ with a zero row wherever $\bar{\alpha}_k = 0$. The columns of $\bar{\boldsymbol{D}}$ correspond to a set of scenarios in which there is a valid set-maximal allocation including at least one losing bid that exactly matches the total $\bar{\alpha}$ amounts of the winning bids. Thus, $\bar{\boldsymbol{D}}$ entries are given by:

$$\bar{\boldsymbol{D}}_{k,\ell} = \begin{cases} 1 & \text{if } x^{\ell}_k = 0 \text{ , } x^{*}_k = 1 \text{ , and } \bar{\alpha}_k > 0 \\ -1 & \text{if } x^{\ell}_k = 1 \text{ , } x^{*}_k = 0 \text{ , and } \bar{\alpha}_k > 0 \\ 0 & \text{otherwise} \end{cases}$$

That is, winning bids not in scenario ℓ get a positive 1 while non-winning bids in ℓ get -1, and the corresponding $\bar{\alpha}$ amounts of these two groups exactly balance for each ℓ column. (Winning bids that are in scenario ℓ effectively cancel out to a zero entry.) Note that the scenarios captured in the columns of $\bar{\boldsymbol{D}}$ are payment-setting to achieve PME by definition. For example, in Example 5, where "any two winners' payments match the combined bids 4 and 5 which total 115," the columns in $\bar{\boldsymbol{D}}$ correspond to the three (symmetric) scenarios in which one winner could join bids 4 and 5 as matching the total payments of all winners.

Under $\bar{\alpha} \neq \mathbf{0}$, at least one equation from (19) is non-zero, and so $\bar{\boldsymbol{D}}$ is not empty. From $\bar{\alpha}\bar{\boldsymbol{D}} = \mathbf{0}$ with nonzero $\bar{\alpha}$, $ker(\bar{\boldsymbol{D}}^T) = \{\alpha | \alpha\bar{\boldsymbol{D}} = \mathbf{0}\}$ is nontrivial, i.e., $nullity(\bar{\boldsymbol{D}}^T) = dim(ker(\bar{\boldsymbol{D}}^T)) \geq 1$, from elementary linear algebra. This *nullity* will be the dimension of our cut space (i.e., the number of cuts produced).

Definition 2. *The set of* critical bids *is* $\bar{\mathcal{K}} = \{k \in \mathcal{K} | \exists \ell \text{ with } \bar{\boldsymbol{D}}_{k,\ell} \neq 0\}$.

Critical bids are those appearing in at least one exact payment-matching scenario, and so must (critically) not see a total payment change during elaboration. Again, assuming an elaboration is needed, this set will be non-empty. The elaboration procedure will generate prices that maintain exactly minimal adjustments $\bar{\alpha}_k$ for all $\bar{\mathcal{K}}$ bids. Notice that $\{k | k \in \mathcal{K}^* \text{ and } \bar{\alpha}_k > 0\} \subset \bar{\mathcal{K}}$ by (19).

Note that not all bids k with $\bar{\alpha}_k > 0$ are in $\bar{\mathcal{K}}$. That is, there are losing bids for which natural linear prices alone cannot satisfy envy-freeness, but which are not strong enough to help form a payment-matching scenario:

Definition 3. *The set of* non-critical bids *is* $\mathcal{K}^{nc} = \{k \in \mathcal{K} | \bar{\alpha}_k > 0 \text{ and } \forall \ell \ \bar{\boldsymbol{D}}_{k,\ell} = 0\}$.

It is not essential (and we find not desirable for interpretability) for non-critical bids to strictly adhere to their minimal adjustment amounts $\bar{\alpha}_k$, and so the elaboration procedure does not enforce

this. (Notice also that there may be a third set of bids, $\mathcal{K} \setminus (\bar{\mathcal{K}} \cup \mathcal{K}^{nc})$, those bids with $\bar{\alpha}_k = 0$. These bids satisfy envy-freeness via natural items alone, and may be referred to as *irrelevant* to the elaboration procedure.)

To form an elaboration of $\bar{\alpha}$ into a set of interpretable cuts, name this set of artificial cuts $\mathcal{J}^+$, with $J^+ = |\mathcal{J}^+|$. These newly-added artificial items (cuts) $\alpha^j \in \mathcal{J}^+$ extend $\mathcal{J}$ (the natural items) and are generated to be concatenated to $\boldsymbol{A}$ as new rows, with prices p_j for $j \in \mathcal{J}^+$ concatenated to p.

Definition 4. *A set of vectors α^j for $j \in \mathcal{J}^+$ is an* elaboration *of $\bar{\alpha}$ if and only if $\exists p \in \mathbb{R}^{J^+}_{\geq 0}$ with:*

$$\sum_{j \in \mathcal{J}^+} p_j \alpha^j_k = \bar{\alpha}_k \qquad \forall k \in \bar{\mathcal{K}} \tag{25}$$

$$\sum_{j \in \mathcal{J}^+} p_j \alpha^j_k \geq \bar{\alpha}_k \qquad \forall k \in \bar{\mathcal{K}}^{nc} \tag{26}$$

$$\alpha^j \boldsymbol{D} \geq \boldsymbol{0} \qquad \forall j \in \mathcal{J}^+ \tag{27}$$

$$\alpha^j \in \mathbb{Z}^K_{\geq 0} \qquad \forall j \in \mathcal{J}^+ \tag{28}$$

Here, (25) define a linear decomposition for critical bids, and (26) relax this decomposition for non-critical bids. The direction of this relaxation ensures envy-freeness is maintained; a non-critical bid will remain envy-free with prices totalling more than $\bar{\alpha}_k$. Critical bids, on the other hand, (those in $\bar{\mathcal{K}}$) match total payments of winning bids or are winning, and so strict equality is required. Constraints (27) assure that each element of the elaboration is a valid cut to WDP, and (28) assures that each artificial cut item consists of non-negative integer amounts in each bundle bid k, emulating natural-item supply constraints (2).

Proposition 8. *For any elaboration, $\alpha^j \bar{\boldsymbol{D}} = \boldsymbol{0}$ for each $j \in \mathcal{J}^+$ with $p_j > 0$.*

Proof. Suppose not: there is some column $\bar{\boldsymbol{D}}^\ell$ with $\alpha^j \bar{\boldsymbol{D}}^\ell > 0$ for some j with $p_j > 0$. (Notice, $\alpha^j \bar{\boldsymbol{D}}^\ell < 0$ is impossible by (27)) Next, right-multiply (25) by $\bar{\boldsymbol{D}}$, resulting in $\sum_{j \in \mathcal{J}^+} p_j \alpha^j \bar{\boldsymbol{D}}^\ell = 0$ for scenario ℓ. This is a contradiction because the LHS of this equation has at least one positive term (by supposition) and no negative terms (by (27) and $p \geq \boldsymbol{0}$). □

Proposition 8 shows that any cut α^j is necessarily in the kernel (aka null space) of $\bar{\boldsymbol{D}}^T$. Intuitively, for any payment-setting scenario to remain payment-setting, changes to the winning bids must match changes to the associated losing bids in each tight scenario. The converse (that a member of the kernel could always be part of some elaboration) is not true in general; first, α^j_k variables with $k \in \mathcal{K}^{nc}$ intersect only zero entries in $\bar{\boldsymbol{D}}$ and so a given kernel basis is orthogonal to (26) constraints; and secondly, constraints (27) may not be satisfied by some elements in the kernel

basis. Thus, while any Gaussian elimination of $\bar{\boldsymbol{D}}^T$ produces a basis for the kernel, members of an arbitrary basis may fail to satisfy (26)-(28) as desired.

We therefore introduce an *elaboration algorithm* with two steps: (i) generate a unique set of basis vectors of $ker(\bar{\boldsymbol{D}}^T)$ that satisfy (25), (27), and (28); and then (ii) use an additional optimization to generate a (perhaps non-unique) lifting of the resulting kernel-basis cuts in order to satisfy the non-critical (26). To accomplish (i) one can iteratively solve a sequence of MIPs and LPs, which we now describe. We find a first kernel-basis cut by solving:

$$\min \quad \alpha^1 \mathbf{1} \tag{Find-α^1}$$

$$\alpha^1 \bar{\boldsymbol{D}} = \mathbf{0} \tag{29}$$

$$\alpha^1 \boldsymbol{D} \geq \mathbf{0} \tag{30}$$

$$\alpha^1 \mathbf{1} \geq 1 \tag{31}$$

$$\alpha^1 \in \mathbb{Z}^K_{\geq 0} \tag{32}$$

Minimizing the sum of the integer α^1 terms reflects our preference for economic interpretability. Constraints assure that the cut is a non-zero (31) integer vector (32) in the kernel (29) and also a valid cut (30).

In order to make progress with each basis cut added, we evaluate how much of $\bar{\alpha}$ can be covered by bases generated thus far by solving an LP, starting at $t = 1$:

$$\min \quad r^t \mathbf{1} \tag{PriceCuts^t}$$

$$\sum_{j=1}^{t} p^t_j \alpha^j + r^t = \bar{\alpha} \tag{33}$$

$$p^t \in \mathbb{R}^t_{\geq 0} \tag{34}$$

$$r^t \in \mathbb{R}^K_{\geq 0} \tag{35}$$

Here r is a vector of remainders, non-negative amounts of $\bar{\alpha}$ that cannot be covered by price-multiples of the t cut vectors generated thus far. When at some iteration t we arrive at zero objective (i.e., $r = \mathbf{0}$), part (i) of elaboration can terminate, having successfully found basis vectors of $ker(\bar{\boldsymbol{D}}^T)$ that satisfy (25), (27), and (28). For example, if $nullity(\bar{\boldsymbol{D}}^T) = 1$, as the lone basis vector of $ker(\bar{\boldsymbol{D}}^T)$, α^1 would generate any kernel element as a multiple, and thus the optimal solution to PriceCuts^t is zero at the first iteration. If not, it is not time to terminate, and so we increment t and generate the next cut vector by solving the following MIP (Find-α^t) for $t \geq 2$:

$$\min \quad \alpha^t \mathbf{1} \tag{Find-α^t}$$

$$\alpha^t + \sum_{j=1}^{t-1} \gamma_j \alpha^j = \delta^t \tag{36}$$

$$\alpha^t \bar{\boldsymbol{D}} = \mathbf{0} \tag{37}$$

$$\alpha^t \boldsymbol{D} \geq \mathbf{0} \tag{38}$$

$$\delta_k \geq 0 \qquad \forall k \in \bar{\mathcal{K}} \mid r_k^{t-1} \geq 0 \tag{39}$$

$$\delta_{\tilde{k}} \geq \epsilon \qquad \text{for some } \tilde{k} \in \bar{\mathcal{K}} \text{ with } r_k^{t-1} > 0 \tag{40}$$

$$\delta_k = 0 \qquad \forall k \in \bar{\mathcal{K}} \mid r_k^{t-1} = 0 \tag{41}$$

$$\alpha^t \in \mathbb{Z}^K_{\geq 0} \tag{42}$$

$$\gamma \in \mathbb{R}^{t-1} \tag{43}$$

$$\delta^t \in \mathbb{R}^K_{\geq 0} \tag{44}$$

As in Find-α^1, each α^t will be an integer vector in the kernel and also a valid cut. Constraints (39) and (40) assure that non-negative progress can be made in reducing positive r^{t-1} terms, with strict progress for an arbitrarily chosen $\tilde{k}$, while not disrupting r^{t-1} terms that have already reached zero. This process of solving Find-α^t then PriceCuts^t continues until $r^t = \mathbf{0}$, at most $t = nullity(\bar{\boldsymbol{D}}^T)$ iterations.

Proposition 9. *After $J^+ = nullity(\bar{\boldsymbol{D}}^T)$ iterations, (Find-α^t, PriceCuts^t) generates artificial-item cuts that satisfy elaboration conditions (25), (27), and (28), and thus form an undominated integer basis of $ker(\bar{\boldsymbol{D}}^T)$.*

Proof. First, Find-α^t is feasible for $t \in \{1, ...J^+\} = \mathcal{J}^+$. The feasible set $\{\alpha \in \mathbb{Z}^K_{\geq 0} | \alpha \bar{\boldsymbol{D}} = \mathbf{0}, \alpha \boldsymbol{D} \geq \mathbf{0}\}$ consists of the integer lattice points of a pointed polyhedral cone, and thus is finitely generated by a unique Hilbert basis (see Aardal et al. 2002). This basis is a collection of component-wise minimal row vectors, together forming a matrix H generating the same feasible set as $\{\alpha | \alpha = pH, p \in \mathbb{Z}^{|H|}_{\geq 0}\}$, which can also be characterized as a Graver basis (see Le and Römer 2024). The theoretical existence of this Hilbert basis guarantees integer solutions to each Find-α^t. Because the Hilbert basis definition demands *integer* weights, $|H| > nullity(\bar{\boldsymbol{D}})$ is possible. Here, these weights are prices and thus continuous, and so a standard linear basis of size $nullity(\bar{\boldsymbol{D}})$ is appropriate. As a consequence, the generated α^t cuts are potentially a strict subset of H: generators of the standard linear cone generated by H. Note that there is no loss in restricting from the integer cone $\{\alpha \in \mathbb{Z}^K_{\geq 0} | \alpha \bar{\boldsymbol{D}} = \mathbf{0}, \alpha \boldsymbol{D} \geq \mathbf{0}\}$ to the linear cone of H, as $\bar{\alpha}$ belongs to the latter, as it is rational given rational coefficients to the original problem, and therefore an integer multiple belongs to the former.

To see that progress in generating bases is made at each t, multiply (36) by an arbitrary $p_t > 0$ and add it to the $t-1$ indexed (33), yielding:

$$p_t \alpha^t + \sum_{j=1}^{t-1} (p_t \gamma_j + p_j^{t-1}) \alpha^j + r^{t-1} - p_t \delta^t = \bar{\alpha} \ .$$

This new α^t extends the basis to improve the objective of PriceCuts^t, because $r^{t-1} - p_t\delta^t \leq r^{t-1}$ by construction, with strict inequality for the chosen $\tilde{k}$. This improvement would not be possible if α^t was a linear combination of previous α^js (or else optimality at $t-1$ is contradicted), thus the basis must grow. Note that this only shows that the new basis establishes an improving direction when p_t is small enough that $p_t\gamma_j + p_j^{t-1} > 0$ for $j < t$, and so we still solve PriceCuts^t in between each Find-α^t iteration. □

The solution set of the series Find-α^t can be solved by an IP solver as written, or be enumerated as in Algorithm 2, presented in Appendix C, which uses a depth-first backtracking search tree. That intuitive characterization can be implemented "by hand" on small instances.

For elaboration part (ii), we need α_k^j for $(j,k) \in \mathcal{J}^+ \times \mathcal{K}^{nc}$ that together satisfy constraints (26). Strengthening the already generated (part (i) elaborated) cuts to include *non-critical* bids is accomplished by solving the following IP, a kind of *group lifting* of these variables:

$$\min \quad \sum_{j\in\mathcal{J}^+}\sum_{k\in\mathcal{K}^{nc}} \alpha_k^j \qquad (nc\text{-lift})$$

$$\sum_{j\in\mathcal{J}^+} p_j\alpha_k^j \geq \bar{\alpha}_k \qquad \forall k \in \mathcal{K}^{nc} \qquad (45)$$

$$\sum_{k\in\mathcal{K}^{nc}} x_k^\ell\alpha_k^j \leq \sum_{k\in\mathcal{K}^*} x_k^\ell\alpha_k^j - \sum_{k\in\mathcal{K}} x_k^\ell\alpha_k^j \qquad \forall (j,\ell) \in \mathcal{J}^+ \times \mathcal{L}^{nc} \qquad (46)$$

$$\alpha_k^j \in \mathbb{Z}_{\geq 0} \qquad \forall (j,k) \in \mathcal{J}^+ \times \mathcal{K}^{nc}$$

Notice that only α_k^j for $(j,k) \in \mathcal{J}^+ \times \mathcal{K}^{nc}$ are decision variables in this formulation, with all p_j, x, and other α_k^j terms fixed from part (i) of the elaboration. Constraints (45) repeat (26), while (46) ensure that each artificial item $j \in \mathcal{J}^+$ remains a valid cut for WDP after lifting in *nc* variables (equivalent to $\alpha\boldsymbol{D} \geq \mathbf{0}$ but focusing the decision variables to the LHS of this formulation and using only $\mathcal{L}^{nc}$ columns of $\boldsymbol{D}$).

To see that this this IP must be feasible, observe that we could selectively increase both bid amount b^k and $\bar{\alpha}_k$ amounts by the same level for some bids $k \in \mathcal{K}^{nc}$ to result in a new auction in which all such k are now critical. In that new auction, elaboration part (i) would give us cuts including $\mathcal{K}^{nc}$, and thus a feasible solution to *nc*-lift. Still, such an adjusted auction approach would tend to overstate the prices quoted to *nc* bids and distort the interpretability of artificial items. We thus pursue the current formulation to find more informative prices.

For the concise examples presented here, we tend not to include extraneous bids that would end up being non-critical. An exception is Example 3: bid b_6 is non-critical because it offers only 34 for all items, when the winning and price setting outcomes are each able to produce at least 35. The elaboration procedure (i) would first generate the cut $x_1 + x_2 + x_3 + x_4 + x_5 \leq 2$ among winning and critical bids, and then (ii) would solve *nc*-lift to lift in b_6 with a coefficient of 2: relative to any

other set-maximal allocation, b_6 knocks out both bids that would pay for one copy of the artificial item. Similarly, if one were to (say) introduce an additional bid $b_6(5A) = 40$ into Example 5, this new bid would be non-critical and lifted in to have coefficients of 1 and 1 for the generated artificial items. This *nc* bid would face a payment of 57.5, which is strictly non-envied and symmetric to the winning bids, which bid on the same set.

As one attempts larger and thus computationally harder auction instances, the treatment of $\mathcal{K}^{nc}$ in elaboration part (ii) could be relaxed or eliminated altogether without abandoning the main results presented. A few choices for elaboration part (ii) are immediately possible:

- **Minimal:** Implement elaboration exactly as stated here. For a selected final outcome based on a particular $\bar{\alpha}$. Non-critical bids are guaranteed to be among the AWE that minimize the total artificial price quotes to *nc* bids.
- **Integer Feasible:** Implement elaboration as stated here, but stopping at any feasible solution to *nc*-lift, perhaps when a computational time-limit or optimality-limit is reached. The *nc* bids see valid linear price quotes that leave them envy-free, but these may not be minimal among such prices, a slightly worse but largely agreeable outcome.
- *nc*-**Nonlinear Only:** Implement elaboration without running part (ii). Final prices are totally explained by the relationship to critical bids, including those that form the price-match allocations. Critical bids see the linear decomposition of $\bar{\alpha}_k = \sum_{j \in \mathcal{J}^+} p_j \alpha_k^j$, as always, while *nc* bids would only see a lump-sum non-linear payment component $\bar{\alpha}_k$ (above linear prices on original items). This lump-sum would cover envy-freeness, and the optimality of Find-$\bar{\alpha}$ would still guarantee that the *nc* bid could not be part of any allocation that would match total payments of the winners.

7. Conclusions and Future Directions

We introduce and examine two main ideas for the computation of prices in a general CA:

- Non-negative integer valid cuts to WDP have natural interpretations as artificial items, and
- PME is a desirable property, with MAP prices as a particular selection of a "best" PME.

Toward the first idea, we introduce matrix $\boldsymbol{D}$ (the directions to optimality from alternate allocations), allowing us to optimize over cuts once a solution to WDP is known. All non-negative integer cuts are in the cone $\alpha \boldsymbol{D} \geq \boldsymbol{0}$ and thus admit a unique Hilbert basis, a set of irreducible "facts" about the relationship of the efficient WDP solution to all other feasible solutions. Pricing these cuts yields the new AWE perspective, and we characterize the closure of cutting algorithms for WDP-LP: any core point can be found as an AWE point by Theorem 1. By Theorem 2 and its corollaries, PME always exist and at one extreme correspond to (i.e., provide a linear decomposition of) MRC payments, a final payment rule which has been used in practice in high-stakes spectrum auctions.

At the other extreme of the PME region, MAP prices emerge as an alternate final-payment rule for any typical CA with the following comparisons to existing payment paradigms:

1. MAP prices provide linear prices while typical MRC and VCG implementations provide only lump-sum payments. (Note: MRC outcomes can always be linearly decomposed to PME with additional AWE cuts, but VCG has such a PME decomposition if and only if VCG = MRC.)
2. By definition, MAP uses the least amount of artificial price adjustment among PMEs, and so is "less artificial" than linearized MRC payments in some cases.
3. Ordered by total payments for a fixed set of bids: pay-as-bid $\geq$ NWE $\geq$ MAP $\geq$ MRC $\geq$ VCG. As intermediate in total-payments, total incentives for misreporting are also intermediate under MAP.
4. If coalitional submodularity holds, min-NWE = MAP = MRC = VCG.
5. Even under the strong substitutes condition, we may have NWE > MAP > MRC.
6. MAP payments are additive over combining unrelated-goods markets, unlike MRC.

As a first paper to explore AWE, PME, and MAP prices, we imagine several avenues for further related research. Given the space limitations of the current submission, details of a larger computational study are deferred to a sequel paper, but we note that MAP requires the solution of several NP-hard WDPs, and so we should expect super-polynomial asymptotic performance unless P = NP. Still, specialized algorithms may prove useful, both for the general problem and for special cases, such as under context-specific preference restrictions.

As with MRC rules (see Bünz et al. 2022), a numerical evaluation of the Bayes-Nash equilibria under MAP pricing would be a complex task in itself. Progress on that difficult computational question could be complemented with behavioral lab experiments or other studies of MAP's robustness to behavioral and structural assumptions, such as relaxing quasi-linearity, risk-neutrality, adding budget constraints, etc. Along a different dimension, we have assumed a single, zero-reserve-normalized seller, which allows us to guarantee core outcomes, and thus guarantee PME existence. Results presented here generalize to a combinatorial *exchange* (with multiple buyers and seller) if one is fortunate enough to have a non-empty core, but when the core is empty, PME is not guaranteed to exist. This opens the possibility of approximate MAP prices when the core is empty, but we leave that for future work.

Finally, the WE concept derives from iterative auctions, where price discovery can serve as a common-value aggregation tool. How best to embed MAP pricing into an iterative price-discovery process remains a promising future direction. Here, we introduce MAP as a sealed-bid format, which could be immediately useful on its own or as the final-payment rule (supplementary round) in a combinatorial clock auction. In future work, we plan to study dynamic MAP pricing, in which AWE cuts are generated dynamically as needed during iterative (clock) rounds.

References

Aardal K, Weismantel R, Wolsey LA (2002) Non-standard approaches to integer programming. *Discrete Applied Mathematics* 123(1):5–74, ISSN 0166-218X, URL http://dx.doi.org/https://doi.org/10.1016/S0166-218X(01)00337-7.

Adomavicius G, Gupta A (2005) Toward comprehensive real-time bidder support in iterative combinatorial auctions. *Information Systems Research* 16(2):169–185.

Alon N, Berman K (1986) Regular hypergraphs, gordon's lemma, steinitz' lemma and invariant theory. *Journal of Combinatorial Theory, Series A* 43(1):91–97, ISSN 0097-3165, URL http://dx.doi.org/https://doi.org/10.1016/0097-3165(86)90026-9.

Ausubel L, Baranov O (2017) A practical guide to the combinatorial clock auction, available at ausubel.com.

Ausubel L, Milgrom P (2002) Ascending auctions with package bidding. *Frontiers of Theoretical Economics.* 1(1).

Ausubel LM, Cramton P, Milgrom P (2006) The clock-proxy auction: A practical combinatorial auction design. *Handbook of spectrum auction design*, 120–140 (Cambridge University Press).

Baldwin E, Bichler M, Fichtl M, Klemperer P (2022) Walrasian equilibria in the strong-substitutes product-mix auction. *Mathematical Programming* 191(2).

Beck MR (2013) *Essays on Core-Selecting Auctions.* Ph.D. thesis, Stanford University.

Bertsimas D, Weismantel R (2005) *Optimization over Integers* (Dynamic Ideas).

Bichler M, Fichtl M, Schwarz G (2021) Walrasian equilibria from an optimization perspective: A guide to the literature. *Naval Research Logistics (NRL)* 68(4):496–513.

Bichler M, Fux V, Goeree J (2018) A matter of equality: Linear pricing in combinatorial exchanges. *Information Systems Research* 29(4):1024–1043.

Bikhchandani S, Mamer JW (1997) Competitive equilibrium in an exchange economy with indivisibilities. *Journal of Economic Theory* 74:385–413.

Budish E (2011) The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119(6):1061–1103.

Bünz B, Lubin B, Seuken S (2022) Designing core-selecting payment rules: A computational search approach. *Information Systems Research* 33:1157–1173.

Conforti M, G C, G Z (2014) *Integer Programming* (Springer).

Day R (2004) *Expressing preferences with price-vector agents in combinatorial auctions.* Ph.D. thesis, University of Maryland.

Day R, Cramton P (2012) Quadratic core-selecting payment rules for combinatorial auctions. *Operations Research* 60:588–603.

Day R, Milgrom P (2008) Core-selecting package auctions. *international Journal of game Theory* 36(3-4):393–407.

Day R, Raghavan S (2007) Fair payments for efficient allocations in public sector combinatorial auctions. *Management science* 53(9):1389–1406.

de Vries S, Schummer J, Vohra RV (2007) On ascending vickrey auctions for heterogeneous objects. *Journal of Economic Theory* 132(1):95–11.

Emadikhiav M, Day R (2024) Walrasian pricing for combinatorial markets with compact-bidding languages: An application to truckload transportation, working Paper.

Gul F, Stachetti E (1999) The english auction with differentiated commodities. *J. of Economic Theory* 92:66–95.

Lahaie S, Lubin B (2019) Adaptive-price combinatorial auctions. *Proceedings of the 2019 ACM Conference on Economics and Computation*, 749–750.

Le D, Römer T (2024) Equivariant lattice bases. *Transactions of the American Mathematical Society* .

Lehmann D, Müller R, Sandholm T (2006) The winner determination problem. *Combinatorial Auctions*, 297–318 (MIT press).

Milgrom P, Watt M (2022) Linear pricing mechanisms for markets without convexity. *Proceedings of the 23rd ACM Conference on Economics and Computation*, 300–300.

Nguyen T, Vohra R (2022) (near) substitute preferences and equilibria with indivisibilities. Technical Report 22-010, Penn Institute for Economic Research, Department of Economics, University of Pennsylvania.

O'Neill RP, Sotkiewicz PM, Hobbs BF, Rothkopf MH, Stewart Jr WR (2005) Efficient market-clearing prices in markets with nonconvexities. *European journal of operational research* 164(1):269–285.

Parkes D, Kalagnanam J, Eso M (2001) Achieving budget-balance with vickrey-based payment schemes in exchanges. *Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI-01)*, 1161–1168.

Appendix A: Dynamic Generation of $\boldsymbol{D}$

Recall from § 6.1 that each column of $\boldsymbol{D}$ is given by $x^* - x^\ell$ for some alternate set-maximal solution x^ℓ of WDP. For small auctions, such as the numbered illustrative examples of this paper, enumerating all such solutions to WDP is easy (not too expensive computationally). As the size of the auction grows, however, the number of solutions grows exponentially, and it becomes practical to only generate columns of the matrix $\boldsymbol{D}$ as needed to guarantee feasibility of an optimization that uses $\boldsymbol{D}$. As is standard in constraint generation, this may be accomplished by solving the *separation problem* associated with the region defined by $\boldsymbol{D}$. Given a partially generated matrix $\boldsymbol{D}^t$ and a potential cut to WDP given by $\alpha x \le \alpha x^*$ (interpreted here as $\alpha_0 = \alpha x^*$ copies of an artificial item) that satisfies $\alpha \boldsymbol{D}^t \ge \boldsymbol{0}$, we have the following:

Proposition EC.1. *Separation at $\boldsymbol{D}^t$: Find solution $x^{\ell+}$ of WDP with bid amounts b replaced by α values. If $\alpha x^{\ell+} \le \alpha x^*$, then $\alpha x \le \alpha x^*$ is a valid cut to WDP. Otherwise, append column $x^* - x^{\ell+}$ to $\boldsymbol{D}^t$ to form $\boldsymbol{D}^{t+1}$ and generate a new potential α.*

Proof. The cut $\alpha x \le \alpha x^*$ is invalid if and only if there is some binary solution x^ℓ to WDP with $\alpha x^\ell > \alpha x^*$. The maximality of $x^{\ell+}$ with objective α establishes the first part. For the second part, notice that the appended column yields $\alpha(x^* - x^{\ell+}) \ge 0$, or $\alpha x^{\ell+} \le \alpha x^*$, so that the same $x^{\ell+}$ will not appear again in a dynamic population of $\boldsymbol{D}$. Note that this process does not require integer α values, and thus can be used in the process of finding the aggregate continuous cut $\bar{\alpha}$. □

In practice, commercial IP solvers like CPLEX and Gurobi offer a solution pool feature, listing high-quality feasible integer solutions found during its internal branching. Feasible solutions thus generated may offer a warm-start to the generation of $\boldsymbol{D}$. Similarly, WDP solutions found in the course of computing VCG payments or MRC payments can also help to pre-populate $\boldsymbol{D}$.

Appendix B: A Linear Mixed-Integer Reformulation of (Find $\bar{\alpha}$)

To compute MAP prices solve:

$$
\begin{aligned}
\min \quad & \bar{\alpha}\mathbf{1} && && \text{(Find } \bar{\alpha}\text{-L)}\\
\text{s.t.} \quad & \bar{\alpha}\boldsymbol{D} \ge \boldsymbol{0} && &&\\
& p\tilde{\boldsymbol{A}} + \bar{\alpha} - \bar{\alpha}^* \ge \tilde{b} && && \text{(EC.1)}\\
& p \le Mz^T && && \text{(EC.2)}\\
& \boldsymbol{A}^{-i}x^{-i} \le c && \forall i \in \mathcal{I}^* &\\
& \boldsymbol{A}^{-i}x^{-i} \ge c - M(\mathbf{1} - z) && \forall i \in \mathcal{I}^* && \text{(EC.3)}\\
& \boldsymbol{B}^{-i}x^{-i} \le \mathbf{1} && \forall i \in \mathcal{I}^* &\\
& \xi^{-i}\mathbf{1} = \bar{\alpha}x^* && \forall i \in \mathcal{I}^* && \text{(EC.4)}\\
& \xi^{-i} \le Mx^{-i} && \forall i \in \mathcal{I}^* && \text{(EC.5)}
\end{aligned}
$$

$$\xi^{-i} \le \bar{\alpha} \qquad \forall i \in \mathcal{I}^* \qquad \text{(EC.6)}$$

$$\xi^{-i} \ge \bar{\alpha} - M(\mathbf{1} - x^{-i}) \qquad \forall i \in \mathcal{I}^* \qquad \text{(EC.7)}$$

$$p\boldsymbol{A}^{-i} + \bar{\alpha} \le b + M(\mathbf{1} - x^{-i}) \qquad \forall i \in \mathcal{I}^* \qquad \text{(EC.8)}$$

$$\bar{\alpha} \in \mathbb{R}^K_{\ge 0}$$

$$p \in \mathbb{R}^J_{\ge 0}$$

$$\xi^{-i} \in \mathbb{R}^K_{\ge 0} \qquad \forall i \in \mathcal{I}^*$$

$$x^{-i} \in \{0,1\}^{K-K_i} \qquad \forall i \in \mathcal{I}^*$$

$$z \in \{0,1\}^J$$

where M is a sufficiently large positive number. Here, (EC.1) substitutes out surplus variables s, resulting in $\tilde{\boldsymbol{A}}$ with each entry given by $\tilde{\boldsymbol{A}}_{j,k} = a^k_j - a^{i(k)*}_j$, each $\tilde{b}_k = b_k - b_{i(k)*}$, and similarly $\bar{\alpha}^*_k = \bar{\alpha}_{i(k)*}$, in order to treat the cut item the same as a natural items in the substitution. Constraints (EC.2) introduce auxiliary binary variables z that must equal 1 when ever a corresponding price p is positive. Then in each sub-economy $\mathcal{I}^* \setminus i$, a positively priced item must not have excess supply according to (EC.3).

Next, continuous ξ^{-i}_k variables replace what had been quadratic terms $\bar{\alpha}_k x^{-i}_k$ in (Find $\bar{\alpha}$). Constraints (EC.4) through (EC.7) respectively enforce that: the fictional item sells fully in each sub-economy; that ξ^{-i}_k can only be positive if the bid k is selected in the sub-economy; and that $\xi^{-i}_k = \bar{\alpha}_k$ exactly when $x^{-i}_k = 1$. Finally, (EC.8) enforces that any bid selected in any sub-economy is willing to cover the price pa^k with respect to IR (and is loose for any bid not selected).

Appendix C: A Search Heuristic for Enumerating Solutions to Find-α^t

Starting with a single starting node, when reaching the Select function, add tree branches from the current node to one new node for each feasible selection k, and proceed forward to a node by selecting an arbitrary k. When a Backtrack command is encountered, run the code backward to the most recent Select command. If there is an unexplored k at that selection, take one such arbitrary k instead of one previously used, and run the code forward with this alternate selection. If there is no unexplored node (because all have been explored on a previous walk forward) then continue backward through the while loop, subtracting previous e_k updates to α, until encountering a Select command with unexplored successors. Proceed forward to explore.

This algorithm is easy to implement "by hand" for small problems and adds intuition for the ease of generating kernel-basis cuts.

Algorithm 2 Enumerate feasible solutions to Find-α^t for $t = 1$ to $t = nullity(\bar{\boldsymbol{D}}^T)$.

Initialize: $keep = \emptyset$; $out = \emptyset$; $\alpha = \mathbf{0}$

Select an arbitrary $k \in \bar{\mathcal{K}} \cap \mathcal{K}^*$, and set $\alpha_k = 1$.

while $\alpha \bar{\boldsymbol{D}} \neq \mathbf{0}$ **do**

Update $\mathcal{K}_{near} = \arg\min_{k \in \bar{\mathcal{K}} | (\alpha + e_k) \notin keep \cup out} \ \|(\alpha + e_k)\bar{\boldsymbol{D}}\|_1$

Update $\mathcal{K}_{best} = \arg\min_{k \in \mathcal{K}_{near}} \ \alpha_k$

Select an arbitrary $k \in \mathcal{K}_{best}$ and set $\alpha \leftarrow \alpha + e_k$

if $\alpha \boldsymbol{D} \geq \mathbf{0}$ **then**

Add α to $keep$

else

Add α to out

if $|keep| = nullity$ **then**

Terminate

else

Backtrack