

New Core-Selecting Payment Rules with Better Fairness and Incentive Properties*

Benjamin Lubin
Information Systems Dept.
Boston University

Benedikt Bünz
Computer Science Dept.
Stanford University

Sven Seuken
Dept. of Informatics
University of Zurich

First version: November 11, 2014

This version: April 28, 2016

Abstract

In this paper, we study the design of core-selecting payment rules for combinatorial auctions, a challenging setting where no strategyproof rules exist. Unfortunately, under the the rule most commonly used in practice, the *Quadratic* rule (Day and Cramton, 2012), large incentives to deviate from truthful bidding remain in equilibrium. Furthermore, the rule significantly favors large over small bidders in terms of incentives as well as expected payoffs. To address this, we propose new payment rules and study them computationally in a new, large domain (with up to 25 goods and 10 multi-minded bidders). We find new rules that outperform the Quadratic rule along all dimensions (efficiency, revenue, incentives, and fairness). Ultimately, we recommend switching to one of our new rules in practice, as it clearly dominates the Quadratic rule.

Keywords: Combinatorial Auctions, Payment Rules, Quadratic, Core, Incentives, Fairness, Gini

JEL: D44 Auctions; D47 Market Design; D82 Asymmetric and Private Information, Mechanism Design

*Emails: blubin@bu.edu, buenz@cs.stanford.edu, seuken@ifi.uzh.ch.

We are thankful for the feedback we received from various participants at the following events (in chronological order): INFORMS'14 (San Francisco, USA), Microeconomics Workshop at HEC-DEEP (Lausanne, Switzerland), and a seminar talk at Microsoft Research (New York, USA). An earlier version of this paper was circulated under the title "Fairness Beyond the Core: New Payment Rules for Combinatorial Auctions."

1. Introduction

The spectrum auctions conducted by governments around the world over the last twenty years are a true success story for market design in general and auction design in particular. Sophisticated mechanisms have been used to sell resources worth billions of dollars, forming the basis for today’s wireless industry (Cramton, 2013). Recent versions of these markets have used a combinatorial auction (CA) mechanism. The advantage of CAs (e.g., in contrast to running multiple single-item auctions) is that buyers can express complex valuation functions over bundles of goods. In this work, we focus on the payments that will be charged after the auction closes and the winners have been determined. Furthermore, we treat the CA as a one-shot game.¹ Thus, we study *direct* payment rules which take as input the bidders’ value reports and compute payments for each of the winning bidders.

At first sight, the famous VCG mechanism may seem like an appealing payment rule for a CA because it is strategyproof. Unfortunately, VCG is generally viewed as unsuitable in a CA domain where items can be complements because it often produces outcomes outside of the *core*. Informally, this means that payments may be so low that a coalition of bidders may be willing to pay more in total than what the seller receives from the current winners. From a revenue perspective, VCG is also often deemed undesirable, because in a CA domain it can produce very low or even zero revenue, despite high competition for the goods in the auction. For this reason, recent auction designs have employed *core-selecting payment rules* (Day and Raghavan, 2007; Day and Milgrom, 2008). The payment rule that is most commonly used in practice, the *Quadratic* rule proposed by Day and Cramton (2012), selects prices that are (1) enforced to be in the core (with respect to the submitted bids), and within this (2) minimal in their total revenue to the seller, and then within this (3) minimal in the Euclidean distance to VCG prices.

Of course, because the Quadratic rule is a core-selecting payment rule, it is not strategyproof, and thus incentives for strategic bidding remain. Only recently has the research community started to grapple with the incentive properties of the Quadratic rule. For the so-called *Local-Local-Global* setting, with 2 items and 3 bidders, Goeree and Lien (2014) as well as Ausubel and Baranov (2013) have independently derived the Bayes-Nash equilibrium of the Quadratic rule. It turns out that, even though the rule minimizes the Euclidean distance to VCG, large incentives to deviate remain in equilibrium. This motivates the search for better payment rules in this paper.

1.1. Bidder Heterogeneity in Real-world Spectrum Auctions

In this paper, we are particularly concerned with studying the “differential” effects of a payment rule on the bidders in an auction, assuming an anonymous auction mechanism (i.e., a mechanism that ignores the bidder identities). If the bidders in the auction were homogeneous (in terms of size, budgets, and values, etc.), then each bidder would, in expectation, have the same incentives, and also receive the same *expected payoff*. However, in real-world auctions, some bidders may be small (e.g., regional) bidders while other bidders may be large (e.g., national) bidders. Naturally, this heterogeneity in size translates into a heterogeneity of budgets, which translates into a heterogeneity of values for bundles of goods in the auction.

For example, an empirical analysis of the bidding data from the British 4G auction in 2013 has identified four heterogeneous classes of bidders (Kroemer, Bichler and Goetzendorf, 2015). The two incumbents, *Vodafone* and *Telefonica* won the most valuable spectrum (800 MHz) that can be used to build out a nation-wide network with maximum reach, and they both paid over £550 Million for their allocations. In contrast, the much smaller bidder *Niche Spectrum Ventures* only won 2.6 GHz

¹We note that the techniques considered are also applicable in the multi-stage designs commonly used in practice.

Bidders	Values			VCG		Quadratic Rule			Fractional Rule		
	A	B	AB	Prices	Payoff	Prices	Payoff	% of VCG Payoff	Prices	Payoff	% of VCG Payoff
1	50*	0	50	10	40	30	20	50%	21.5	28.5	~71%
2	0	150*	150	50	100	70	80	80%	78.5	71.5	~71%
3	10	0	100	-	-	-	-	-	-	-	-

Table 1: Motivating Example: Comparing VCG, the Quadratic Rule, and the Fractional Rule. The * sign indicates the bids selected under the efficient allocation.

spectrum, whose reach is unsuited to build out a nation-wide network, and paid roughly £200 Million for its allocation. Other, even smaller bidders, submitted bids (but did not win) suggesting a valuation of roughly £10 Million for the least valuable spectrum band.²

Another example is the FCC Auction-73 from 2008, which generated approximately \$19 Billion in revenue. Out of the 214 qualified bidders, 101 were winners.³ Their winning bid amounts ranged between \$20,000 for one of the smallest bidders to more than \$9 Billion (Verizon Wireless).⁴ Finally, the most recent FCC Auction-97, which closed in January 2015 and generated over \$40 Billion in revenue, had 31 winning bidders. Their total winning bid amounts ranged between \$155,000 for one of the smallest bidders, to \$18 Billion for AT&T and \$10 Billion for Verizon.⁵

Clearly, we observe a large bidder heterogeneity in spectrum auctions conducted in practice. This raises the following questions: (1) Do some auction payment rules lead to an advantage for large bidders, while other rules lead to an advantage for small bidders? (2) What is the impact on efficiency and revenue when small or large bidders are being favored? (3) Given all of these considerations, what choice should a government make when selecting a particular payment rule for a government-run combinatorial auction? This leads us to study new aspects in the design of payment rules, in particular the dimensions of “aggregate incentives” as well as “fairness.”

1.2. Fairness: A Motivating Example

So far, we have focused our discussion on comparing the incentive properties of the Quadratic rule with other rules (e.g. the Small rule). In this section, we now provide a motivating example, to build some intuition for the “unfairness” of the Quadratic rule in terms of the *payoff distribution* it produces. Note that for the following example only (for illustrative purposes), we assume that the bidders submit their true values, and do not play their BNE strategies. All analyses in the rest of the paper will be in BNE.

Consider Table 1, which contains one row for each of the bidders 1, 2 and 3.⁶ The efficient allocation gives item *A* to bidder 1 and item *B* to bidder 2. VCG prices are $p_1^{VCG} = 10$ and $p_2^{VCG} = 50$, with payoffs of $\pi_1^{VCG} = 40$ and $\pi_2^{VCG} = 100$. We immediately see that VCG is outside the core (Day and Raghavan, 2007): bidders 1 and 2 are only paying 60 in total for the two items *A* and *B*, which is less than the 100 offered by bidder 3, and thus a violation of the core. Clearly, this is *unfair* towards bidder 3 who is offering 100 for items that are sold for only 60.

²<http://stakeholders.ofcom.org.uk/binaries/spectrum/spectrum-awards/awards-in-progress/notices/4g-final-results.pdf>

³http://wireless.fcc.gov/auctions/default.htm?job=auction_summary&id=73

⁴https://apps.fcc.gov/edocs_public/attachmatch/DA-08-595A3.pdf

⁵https://apps.fcc.gov/edocs_public/attachmatch/DA-15-131A3.pdf

⁶To simplify the language, we will use “he” when referring to a bidder, and “she” when referring to the seller.

The Quadratic Rule, presented next in Table 1, corrects this unfairness present in VCG by choosing a price vector from the core, i.e. it picks a price vector with revenue at least 100, which is bidder 3’s offer. Among such vectors the Quadratic Rule selects the one which is minimal in revenue and closest by the Euclidean distance to the VCG prices, i.e., in this example $p_1^Q = 30$ and $p_2^Q = 70$. Because the payment of both bidder 1 and 2 increased over that in VCG, their payoffs decreased to $\pi_1^Q = 20$ and $\pi_2^Q = 80$. Note, however, that this payoff decrease has impacted the two bidders unequally. Bidder 1 is left with only 50% of his VCG payoff, while bidder 2 retains 80% of his VCG payoff. This effect is expected under the Quadratic rule. Given a single core constraint, the rule will increase the prices uniformly, starting from VCG, until the constraint is met. Naturally, this will decrease the relative payoffs, i.e. the payoffs as a fraction of VCG, of bidders with smaller VCG payoffs by more than those with larger payoffs. This leads to the observed unfairness of the Quadratic rule in terms of the relative payoffs of the winners, here between the *small* bidder (i.e., bidder 1) and the *large* bidder (i.e., bidder 2).

As we have seen, the Quadratic rule corrects the unfairness of VCG, by selecting payments in the core (an idea put forward by (Day and Raghavan, 2007)), but still leads to unequally distributed relative payoffs among the winners. In this paper, we offer a new fairness notion that goes *beyond the core* to account for this. An intuition for this is provided by the *Fractional* rule in Table 1, with $p_1^{MF} = 21.5$ and $p_2^{MF} = 78.5$. Note that these prices still produce a revenue of 100 and are in the core. Yet, both bidders get the *same fraction of their VCG payoff* (roughly 71%). Thus, this rule has both the fairness of the core and the fairness of uniform relative payoffs at the same time. All else equal, we argue that it is therefore fairer than the Quadratic rule. But of course, our goal will be to find a payment rule that achieves this, not given truthful value reports by the bidders, but in BNE.

1.3. Measuring Fairness in Payoff Space via the Gini Coefficient

As we have seen in the motivating example presented in Table 1, the three different payment rules (VCG, Quadratic, Fractional) produce different payoff distributions. But how do we now formally compare the payment rules regarding how *fair* they are? Obviously, we need a measure for the degree of fairness of a payment rule.

In this paper, we offer two such measures. We first consider a very basic, straight-forward measure of fairness: the *standard Gini* coefficient, as it is commonly used in Economics to compare the degree of (in-)equality of income within a country, for example. Concretely, we compute the Gini coefficient over the expected payoffs of all bidders. Computing this Gini coefficient can be done analytically (e.g., for small settings like the LLG domain) as well as numerically (for larger settings). The Gini coefficient provides us with a number between 0 and 100 for each payment rule, where a smaller number indicates a “fairer” payment rule. Thus, using this approach, we can compare the fairness of all payment rules that we consider in this paper.

One problem with applying the standard Gini in an auction domain with heterogeneous bidders is that we actually expect larger bidders (with larger values, who probably contribute more to social welfare) to obtain larger payoffs than smaller bidders. Such an *unequal* distribution of payoffs is absolutely fine and to be expected, in an auction domain with small and large bidders (in contrast to the income distribution within a country, which some may argue should not be too unequal). Thus, the standard gini may not be the only goal we should consider when designing a core-selecting payment rule.

For this reason, we have also constructed a second fairness measure that is geared specifically towards measuring the fairness properties of payment rules in auctions with heterogeneous bidder populations: a *VCG-based Gini coefficient*. We argue that it is fine for differently-sized bidders to expect different

payoffs, and that those should ideally be equal to their contribution to social welfare: a “maximally fair” payment rule would provide every bidder with payoffs equal to his social welfare. Of course, VCG is the unique mechanism that achieves this. More precisely, VCG is the unique mechanism with the following two properties:

1. Every bidder’s *payoff* is his *contribution* to the social welfare.
2. Every bidder’s *payment* is the *externality* he imposes on all other bidders.

Of course, when payments are required to be in the core, we are often not able to charge VCG payments anymore and thus bidders will not receive their VCG payoffs. In particular, because core revenues are generally higher than VCG revenues, the bidders’ payoffs under a core-selecting mechanism will generally be lower than under VCG. However, a core-selecting payment rule can still provide each bidder with a certain *fraction of his VCG payoff*, with bidders obtaining a larger fraction of their payoff consequently receiving better treatment.⁷

Using these fractions of VCG payoff, we adopt as our main fairness measure in this paper a *perfectly balanced distribution of VCG payoff fractions across all bidders* in the auction, as measured by a Gini coefficient (we detail our precise method for calculating a Gini coefficient for payment rules in Section 4.3). Note that the *Fractional* rule in Table 1 had indeed maximal fairness according to this fairness measure (given truthful value reports), while the Quadratic rule did not. Of course, the VCG mechanism always has perfect fairness, but in general produces payments that are outside the core.

In Section 7, where we evaluate our payment rules in Bayes-Nash equilibrium, we always provide the standard Gini as well as the VCG-based Gini. However, we will see that, surprisingly, the two measures lead to the same ranking of our rules in terms of fairness. In particular, across all domains that we study in this paper, the Small and Fractional rules (with various weightings and reference points) are much fairer than the Quadratic rule, according to both measures of fairness.

1.4. Comparing Payment Rules in Realistically-sized Combinatorial Auction Domains

As mentioned before, in this paper we develop a new framework which enables us to systematically evaluate 86 distinct payment rules for combinatorial auctions. This, of course, immediately brings up the question of what the right “evaluation method” should be. Because no strategyproof core-selecting payment rules exist, we must evaluate our payment rules *in equilibrium*. Much of the early work on core-selecting auctions has used a *full information Nash equilibrium* analysis. However, we argue that the full information assumption is simply too unrealistic for practical settings, in particular for large combinatorial auctions with many heterogeneous bidders. For this reason we adopt a *Bayes-Nash equilibrium* (BNE) analysis to evaluate and compare our payment rules.

We are interested in understanding the performance of our rules in realistically-sized auction domains. Unfortunately, analytically deriving the BNE quickly becomes intractable when going beyond the LLG setting.

For this reason, we have developed and implemented a computational approach to find an *approximate Bayes-Nash equilibrium* for our payment rules. More concretely, we have developed an algorithm, based on fictitious play, which takes as input a payment rule and a domain generator (i.e., a simulator for a combinatorial auction domain), and then searches for an ϵ -BNE. Using this approach, we can find an approximate BNE for one payment rule for a domain with 25 items and 10 multi-minded bidders on a high-performance machine within approximately 12 hours. Of course, we used a compute cluster to parallelize this computation for all the payment rules and domains we have studied.

⁷We view a relative measure between provided payoff and VCG payoff as appropriate because it accounts for the large heterogeneity in the underlying payoffs available and corresponding market power, which an absolute scale would not.

1.5. Overview of Contributions

Our ultimate contribution in this paper is a new payment rule, the *Marginal-Economy-Fractional* rule, which outperforms the Quadratic rule, and strikes the best balance along all of our performance criteria. To get there, we offer:

1. A new *framework* for the design of core-selecting payment rules, and two *fairness measures based on the Gini coefficient* for evaluating payment rules.
2. The design of 75 new rules that have not been studied before, built around various (i) distance metrics, (ii) reference points, and (iii) bidder weightings.
3. A set of novel *Marginal Economy* weights that perform almost as well as reserve price weights but without requiring *a priori* knowledge on the part of the seller.
4. An algorithm for computationally finding approximate BNEs in realistically-sized core-selecting CAs.
5. Extensive experimental results using our computational BNE approach, comparing 85 different rules in terms of efficiency, aggregate incentives, fairness, and revenue.

We find that, in BNE, the Quadratic rule is outperformed by many of our new rules in terms of aggregate incentives and fairness, while efficiency is generally high for all rules. Additionally, we show that the minimum-revenue-core (MRC) constraint, typically included in today’s core-selecting CAs, only has a minor effect in equilibrium. After reviewing all results, we find that one rule is the best all-rounder: the *Marginal-Economy-Fractional* rule with VCG as the reference point. We find that this rule generally performs as well as or even slightly better than the Quadratic rule in terms of efficiency and revenue, but significantly improves upon it in terms of aggregate incentives and fairness.

2. Related Work

Over the last 15 years, there has been a large literature on the design of bidding languages, clearing algorithms, and activity rules for use in combinatorial auctions (CAs) (Cramton, Shoham and Steinberg, 2006). But finding optimal *payment rules*, the focus of our paper, has turned out to be somewhat elusive. Early proposals for CAs typically considered VCG prices (Varian and MacKie-Mason, 1994). However, as has been pointed out in the literature (Ausubel and Milgrom, 2006), there are numerous issues with VCG, most notably that it may result in arbitrarily low revenue to the seller. This observation has led to considerable interest in alternative payment rules in recent years, and ultimately the proposal of *core-selecting combinatorial auctions* (Ausubel and Milgrom, 2002; Milgrom, 2007; Day and Milgrom, 2008).

Unfortunately, there exists no strategyproof core-selecting payment rule, and thus, designing an “optimal” core-selecting rule is a challenging market design problem. Othman and Sandholm (2010) avoid this issue by assuming participants bid so as to minimize envy instead of maximize profit. Instead, we follow the line of research that has made the more common assumption that bidders seek to maximize profit. In this direction, Parkes, Kalagnanam and Eso (2001) were the first to introduce the idea of finding prices that minimize some distance metric to VCG, first for combinatorial exchanges, and later for CAs (Parkes, 2002). Since then, a small number of CA payment rules have been proposed that minimize some distance metric to VCG (Day and Raghavan, 2007), or to other reference points (Erdil and Klempere, 2010). Ultimately, Day and Cramton (2012) proposed the *Quadratic* rule, and

this is also the rule most often used in practice today. We will describe all of the payment rules and explain the origins of the underlying ideas in detail in Section 5.6.

Even though the Quadratic rule has now been used for about 5 years by multiple governments to allocate resources worth Billions of dollars, we still have an incomplete understanding of this rule. For example, the only Bayes-Nash equilibrium results we have for the Quadratic rule are for highly stylized settings with two items and three bidders (Goeree and Lien, 2014; Ausubel and Baranov, 2013). To study core-selecting rules in larger settings, where it is infeasible to obtain analytical BNEs, we compute approximate BNEs, following earlier work by Lubin and Parkes (2009) who have evaluated pricing rules for combinatorial *exchanges* via computational BNE approximations. Our computational techniques build on this approach, as well as earlier work from Reeves and Wellman (2004) and Vorobeychik and Wellman (2008).

3. Preliminaries

3.1. Formal Model

In a *combinatorial auction (CA)*, there is a set M of m distinct, indivisible items, and a set N of n bidders. Each bidder i has a *valuation function* v_i which, for every bundle of items $S \subseteq M$, defines bidder i 's value $v_i(S) \in \mathbb{R}$, i.e., the maximum amount that bidder i would be willing to pay for S . To simplify notation, we assume that the seller has zero value for all items, although our setup extends to the case where the seller has non-zero value (see Day and Cramton (2012) for how to handle reserve prices).

We let $p = (p_1, \dots, p_n)$ denote the payment vector, with p_i denoting bidder i 's payment. We assume that bidders have quasi-linear utility functions, i.e., $u_i(S, p_i) = v_i(S) - p_i$. Bidders make reports about their values to the mechanism, denoted $\hat{v}_i(S)$, which may be non-truthful (i.e., $\hat{v}_i \neq v_i$). Following existing work in this area (Goeree and Lien, 2014; Ausubel and Baranov, 2013), we assume that bidders only bid on items for which they have a positive value. We define an *allocation* $X = (X_1, \dots, X_n) \subseteq M^n$ as a vector of bundles, with $X_i \subseteq M$ being the bundle that i gets allocated. A mechanism's *allocation rule* maps the bidders' reports to an allocation. We only consider allocation rules that maximize reported *social welfare*, yielding an allocation $X^* = \arg \max_X \sum_{i \in N} \hat{v}_i(X_i)$, subject to X being feasible, i.e., $\bigcap X_i^* = \emptyset$. In addition to the allocation rule, a mechanism also specifies a *payment rule*, defining prices. Together, these define the *outcome* $O = \langle X, p \rangle$. An outcome O is called *individually rational (IR)* if, $\forall i: u_i(X_i, p_i) \geq 0$.

3.2. VCG, Payoff, and Discount

The famed VCG mechanism (Vickrey, 1961; Clarke, 1971; Groves, 1973) generalizes the well-known second-price auction:

Definition 1 (VCG). *The VCG mechanism chooses the allocation X^* that maximizes the bidders' reported social welfare, and then charges each bidder i its marginal cost to the economy, i.e., the externality he imposes on all other bidders. Formally:*

$$p_i^{VCG} = \sum_{j \neq i} \hat{v}_j(X^{-i}) - \sum_{j \neq i} \hat{v}_j(X^*), \quad (1)$$

where X^{-i} is the welfare-maximizing allocation when all bidders except i are present.




Figure 1: A graphical depiction of the core, as well as the price vectors corresponding to VCG, Quadratic, Δ -Large, Δ -Small, Δ -Fractional, and Pay-as-Bid.

VCG is social-welfare maximizing and *strategyproof*, i.e., it is a dominant strategy for every bidder to report his true value v_i . Every payment rule other than VCG we consider in this paper will not be strategyproof, and thus we will in general have to differentiate between a bidder’s true value v_i and his reported value \hat{v}_i .

Definition 2 (Payoff). *Given a payment p_i charged by a particular payment rule, a bidder’s payoff is the bidder’s profit evaluated at his true value, i.e.,*

$$\pi_i = v_i(X_i^*) - p_i. \quad (2)$$

Definition 3 (Discount). *A bidder’s reported VCG discount (or just discount) is the difference between its reported value and its VCG payment:*

$$\Delta_i = \hat{v}_i(X_i^*) - p_i^{VCG}. \quad (3)$$

Note that *payoff* and *discount* are only the same (in equilibrium) if VCG is used as the actual payment rule; otherwise they are different. The *payoff* is what a bidder actually cares about, and the *discount* captures the “revealed payoff” the bidder *would have gotten* if VCG had *instead been* used in place of the actual payment rule. We will later use the *discount* for the design of our payment rules to weight the bidders’s bids.

3.3. The Core

In CAs where some items are complements, VCG prices may lie outside the *core*. Informally, this means that a coalition of bidders is willing to pay more than what the seller receives from the current winners. To avoid such undesirable outcomes, recent designs have employed payment rules that restrict prices to be in the *core*, giving rise to *core-selecting payment rules* (Day and Milgrom, 2008).

Definition 4 (Core Prices). *We let W denote the set of winners, X^* the welfare-maximizing allocation, $C \subseteq N$ denotes a coalition of bidders, and X^C is the allocation that would be chosen by the mechanism if only the bidders in the coalition C were be present. Then, a price vector p is in the core, if, in addition to individual rationality, the following set of core constraints hold:*

$$\sum_{i \in W} p_i \geq \sum_{i \in C} v_i(X^C) - \left(\sum_{i \in C} v_i(X^*) - \sum_{i \in C} p_i \right) \quad \forall C \subseteq N \quad (4)$$

Note that enforcing prices to be in the core puts lower bounds (constraints) on the payments of the winners, where each coalition of bidders leads to one core constraint. Intuitively, the winners' payments must be *sufficiently large*, such that there exists no coalition that is willing to pay more to the seller than the current winners' payments. In a CA with complements, VCG prices are often outside the core and, in the worst case, VCG may generate zero revenue despite high competition for the goods. See Figure 1 for a graphical depiction of the core, based on the motivating example from Section 1.2.

In this example, we immediately see that VCG prices are outside the core. Additionally, we also display the price vectors that would be chosen by the Quadratic rule (which we describe next), as well as the Large, Small, and Fractional rules (which we describe in Sections 3.5 and 5.4). Note that the distance between VCG and the core can vary significantly depending on the bidders' bids (i.e., from one auction to another). In our experiments, we found that this distance has a significant impact on how fair different payment rules are, and, not surprisingly, how much revenue they generate (see Section 7.4).

Remark 1 (True vs. Revealed Core). *Note that the core is defined in terms of bidders' true values. However, given that no strategyproof core-selecting CA exists, we must expect bidders to be non-truthful. Goeree and Lien (2014) have recently shown that the outcome of a core-selecting CA can be outside the true core in a BNE. Thus, core-selecting CAs only guarantee to produce outcomes in the revealed core, i.e., in the core with respect to reported values. Going forward, whenever we talk about the core, or core-selecting rules, we always mean the revealed core, unless we state it otherwise. Note that for the auctioneer (e.g., the government), having prices in the revealed core is typically what she wants, because this protects the auctioneer against law-suits from losing bidders and from the appearance that an unacceptably large amount of revenue was left on the table.*

3.4. Bayes-Nash Equilibrium

For the analysis of auction payment rules we assume that the bidders know their own value function, but do not have full information about the other bidders' value functions. Instead, bidders only have distributional information regarding the other bidders' value functions. Thus, the appropriate equilibrium concept is the *Bayes-Nash Equilibrium*. For the following definition, we let s_i denote bidder i 's strategy, which is a mapping from his true value function v_i to a possibly non-truthful report \hat{v}_i . Given a value function and a strategy from each bidder, this determines the outcome of the auction. Thus, we can let $u_i(s_1(v_1), s_2(v_2), \dots, s_n(v_n))$ denote bidder i 's utility for the outcome of the auction. We use v_{-i} to denote the value functions of all bidders except i , and analogously for the strategies s_{-i} :

Definition 5 (Bayes-Nash Equilibrium). *A strategy profile $s^* = (s_1^*, \dots, s_n^*)$ is a Bayes-Nash equilibrium (BNE) in a sealed-bid auction if, for all bidders i , and all values functions v_i*

$$\mathbb{E}_{v_{-i}} [u_i(s_i^*(v_i), s_{-i}^*(v_{-i}))] \geq \mathbb{E}_{v_{-i}} [u_i(\hat{v}_i, s_{-i}^*(v_{-i}))], \quad \text{for all possible reports } \hat{v}_i, \quad (5)$$

where the expectation is taken with respect to the distribution over the other bidders' value functions.

In words, a bidder's BNE strategy is an optimal strategy (a mapping of true value functions to reported value functions) given his belief regarding the other bidders' value functions. With this background, we turn to describing the key desiderata for payment rules in the core-selecting CA setting.

3.5. Existing Theoretical Analysis of the Quadratic Rule

To place our subsequent results in context, we here include the existing analytical results for the Quadratic rule in the tractably-small LLG domain (Ausubel and Baranov, 2013; Goeree and Lien, 2014), which we first describe.

3.5.1. The LLG Setting

Bayes-Nash equilibria of core-selecting payment rules are very complex to study analytically, which is why theoretical results only exist for a very small settings. Existing theoretical results have been derived for the *Local-Local-Global* (LLG) setting. In this setting, there exist two items A and B , and three bidders. Two of the bidders are local bidders, with one of them only interested in item A , and one of them only interested in item B . The third bidder is the global bidder who is interested in both items A and B . In this set-up the global bidder has a weakly dominant strategy to bid truthful. However, the local bidders have an incentive to “shade” their bids.

3.5.2. The Quadratic Rule

The *Quadratic* rule is the payment rule that has been used in many of the spectrum auctions in North America and Europe over the past five years (Day and Cramton, 2012):

Definition 6 (Quadratic Rule). *The unique Quadratic Rule price vector is chosen to be:*

1. *Within the core (defined with respect to reported values).*
2. *Within this, minimal in the L_1 (Manhattan) norm, the so called minimal revenue core (MRC), an idea first put forward by Day and Raghavan (2007).*
3. *Within this, minimal in the L_2 (Euclidean) distance to the VCG payments,⁸ an idea first put forward by Day and Cramton (2012).*

The BNE for the specific set-up where the value of the global bidder is uniformly distributed on $[0,2]$ and the value of the local bidders are uniformly distributed on $[0,1]$ and where all such values are drawn independently from each other and are uncorrelated, has already been calculated:

Proposition 1 (Goeree and Lien (2014); Ausubel and Baranov (2013)). *The Bayes-Nash equilibrium of the Quadratic rule with uncorrelated bids is for the global bidder to bid truthful, and for the local bidders to bid:*

$$\hat{v} = \max(0, v - (3 - 2\sqrt{2})) \approx \max(0, v - 0.17) \quad (6)$$

Thus, the Bayes-Nash equilibrium strategies of the local bidders require an *additive* shading of their values in this uncorrelated LLG setting.

The case of correlated bidders has also already been studied:

Proposition 2 (Ausubel and Baranov (2013)). *The Bayes-Nash equilibrium of the Quadratic rule with perfectly correlated local bidders is for the global bidder to bid truthful, and for the local bidders to bid:*

$$\hat{v} = \frac{2}{3}v \quad (7)$$

Thus, it turns out that in this setting, the Bayes-Nash equilibrium strategies of the local bidders require a *multiplicative* shading of their values. Ausubel and Baranov (2013) also derive Bayes-Nash equilibrium strategies for LLG settings with partial correlation which will induce the local bidders to shade both additively and multiplicatively according to a more complicated formula we omit here for brevity.

⁸Note that the L_1 restriction from step (2) can be binding. Thus, this rule is not the same as simply minimizing the L_2 distance to VCG without first minimizing L_1 .

4. Design Goals

We strive for the following four objectives when designing core-selecting payment rules, in this order: (1) high efficiency, (2) good incentives, (3) good fairness properties, and (4) high revenue. Obviously, trade-offs are necessary.

4.1. High Efficiency

From a social planner’s perspective (e.g., a government auctioning off wireless spectrum) it is desirable to maximize social welfare, i.e., *efficiency*. Of course, the efficiency of our payment rules must be evaluated in BNE, and thus we consider the *expected* efficiency in BNE. Our measure for efficiency is the expected welfare of the mechanism divided by the expected welfare of the optimal allocation. Formally, given an auction instance I , let $SW_{OPT}(I)$ denote the social welfare obtained under the optimal allocation given bidders’ true values (which is equal to the social welfare that would be achieved by VCG). Let $SW_M(I)$ denote the social welfare obtained by the mechanism M when all bidders play their BNE strategies. We define the efficiency of mechanism M as:

$$\text{Efficiency}(M) = \frac{\mathbb{E}_{\sim I}[SW_M(I)]}{\mathbb{E}_{\sim I}[SW_{OPT}(I)]} \quad (8)$$

where the expectation is taken over the auction instances given a particular domain generator. This is the standard definition of efficiency used in prior work, e.g., by [Goeree and Lien \(2014\)](#).⁹ Of course, we could simply use VCG to achieve the maximally efficient allocation. However, as discussed in the introduction, VCG outcomes are often outside the core, and finding payments inside the core is also one of our design goals.¹⁰ Thus, maximizing efficiency in BNE is our first non-trivial design goal.

4.2. Good Aggregate Incentives

We next desire that our payment rules produce “good incentives”. Because there is no strategyproof core-selecting CA, and thus, all of our rules are necessarily non-strategyproof, there will remain strategic opportunities for the participants. However, we would like these opportunities in BNE to be as small as possible. This helps unsophisticated players who play truthfully instead of playing optimally, because they will suffer less if their BNE strategy is closer to playing truthful, and in general makes the game easier for participants to play. Furthermore, from an “approximate strategyproofness” point of view, smaller opportunities to manipulate in BNE are also better: one can take the viewpoint that an agent who can only gain very little (in BNE) may not want to manipulate at all.

To make this notion of “good aggregate incentives” concrete, we define the *aggregate incentives* for bidders to strategize as the Euclidean distance between their truthful value and their bid in BNE, normalized by their truthful value. The normalization will make the measure robust to the distribution of value in the domain. Formally we have:

⁹Note that [Ausubel and Baranov \(2013\)](#) also study core-selecting CAs in BNE, but what they call *efficiency* is actually not equivalent to our definition. Instead, their concept of efficiency corresponds to the *probability of achieving the efficient allocation*, which is an important difference. They find that for the Quadratic rule, in some settings, the probability of achieving the efficient allocation is only 85% in BNE. However, in that same setting, the social welfare of the mechanism (in terms of the sum of bidders’ values) is still above 98%.

¹⁰Note that we put “efficiency” first, and “payments in the core” second, in our list of objectives, because if requiring payments in the core would lead to a large drop in efficiency, we would probably prefer a payment rule that has larger efficiency but is not core-selecting over a payment rule that is highly inefficient (in BNE) but is core-selecting. Fortunately, as our computational BNE analyses will show, all of the core-selecting payment rules we consider achieve near 100%-efficiency, and thus, requiring payments in the core does not lead to a significant loss in efficiency.

Definition 7 (Aggregate Incentives).

$$AI = \|(v-\hat{v})/v\|_2, \quad (9)$$

where the L_2 measure is taken over a vector calculated from the observed and reported values.

We note that in Section 3.5, we will use a continuous version of this measure; later in Section 7, where bidders are organized by class, we will use a discretized version.

4.3. Fairness: Gini Coefficients for Payment Rules

As we have motivated in Section 1.2, the core only provides a minimal fairness guarantee (“no-group-envy”), but says little about the distribution of winners’ payoffs. In this section, we define two measures for how fair the payoff distribution of a payment rule is. While any measure of statistical dispersion might be employed, we use the *Gini coefficient*, which is often used for this purpose in the economics literature. Accordingly, we define a “Standard” Gini measure over these payoffs as:

Definition 8 (Standard Gini). *Given a vector of participant payoffs π^R , we define the Standard Gini as the unweighted Gini coefficient computed on these values (Jasso, 1979). Here π_i^R refers to agent i ’s payoff under rule R in the (approximate) BNE of rule R .*

Gini coefficients are between 0 and 100: a Gini of 0 implies perfect fairness and a Gini of 100 implies perfect unfairness over the quantity of comparison.

However, because the Gini is generally used to measure wealth inequality, some care must be used in applying it to our more complex setting. We believe the payoff obtained under the rule to be the most natural analogy to “wealth” in our setting, but it is certainly not the only quantity we might wish to consider. Specifically, as we have argued in Section 1.3, this is not the only measure of fairness that we may wish to consider. We may also reasonably wish to adopt VCG as our gold standard in terms of fairness, even as it is one that is not achievable within the core. We might then wish to determine to what degree our core-selecting payment rules allocate to participants the same *fraction of their VCG payoff*, and use this as our fairness target:

Definition 9 (Fraction of VCG Payoff). *Given bidder i ’s VCG payoff π_i^{VCG} and his payoff under rule R , i.e., π_i^R , we define bidder i ’s fraction of VCG payoff as:*

$$\gamma_i = \frac{\pi_i^R}{\pi_i^{VCG}} \quad (10)$$

where π_i^{VCG} refers to agent i ’s VCG payoff given truthful value reports (the BNE of VCG), while π_i^R refers to agent i ’s payoff under rule R in the (approximate) BNE of rule R .

Accordingly, we will now quantify the degree to which a payment rule achieves this notion of fairness by seeking a balanced distribution of VCG payoff fractions across the bidders. To begin, we use the standard measure to capture the fairness of the winning bidders as follows:

Definition 10 (Winners’ Gini). *Given a vector of the fractions of VCG payoffs, i.e., $\gamma = (\gamma_1, \dots, \gamma_n)$, we define the Winners’ Gini for γ as the Gini coefficient computed on the set of bidders who are winners under both rules, R and VCG:*

$$\text{Winners’ Gini}(\gamma) = G(\gamma_i : i \in W^R \cap W^{VCG}) \quad (11)$$

where G is the standard unweighted Gini coefficient (Jasso, 1979) and W^R and W^{VCG} are the sets of winners under R and VCG respectively.

In practice, we draw a large number of bidders i from our domain of interest, i.e. we run a vary large number of auctions (100,000) to produce γ .

Note, however, that the Winners' Gini covers only those bidders who are winning under both, the rule R and VCG. To get a handle on what fraction of bidders this represents, we also consider the following:

Definition 11 (Winner Change Rate). *Given W^{VCG} , i.e., the set of winners under VCG, and W^R , i.e., the set of winners under R , we define the Winner Change Rate of R :*

$$\text{Winner Change Rate}(R) = 1 - \frac{|W^{VCG} \cap W^R|}{|W^{VCG} \cup W^R|} \quad (12)$$

This is the fraction of bidders who are winners under one rule but not the other, otherwise known as as the Jaccard distance $d_J(W^{VCG}, W^R)$.

The bidders captured by the Winner Change Rate are precisely those that create a maximal *unfairness*, because their fraction of VCG payoff is either 0 or ∞ .

Ultimately, we can combine the *Winners' Gini* and the *Winner Change Rate* into a single fairness measure for all bidders:

Definition 12 (VCG Gini). *Given a vector of the fractions of VCG payoffs, i.e., $\gamma = (\gamma_1, \dots, \gamma_n)$, we define the VCG Gini for γ as:*

$$\text{VCG Gini}(\gamma) = \text{Winners' Gini}(\gamma) \cdot (1 - \text{Winner Change Rate}) + 100 \cdot \text{Winner Change Rate}$$

The VCG Gini captures the Winners' Gini for those bidders in $W^R \cap W^{VCG}$, and assigns a Gini of 100 explicitly to those bidders in $\{W^R \cup W^{VCG}\} \setminus \{W^R \cap W^{VCG}\}$.

The VCG Gini provides a single fairness number for a payment rule R that captures both, the change in the winners' between VCG and the rule R , as well as the distribution of the fractions of VCG payoffs among the winners under R .

4.4. High revenue

When choosing between different rules, various design goals are possible, including the natural goal of *maximizing revenue*. However, in this paper we are primarily interested in *fairness-maximizing payment rules*, and not revenue-maximizing rules. If the seller wanted to achieve especially high revenue, she would have to use well-chosen reserve prices anyway, which of course would require distributional information about bidders' values. Because this is not the focus of this paper, we do not consider using reserve prices to increase revenue in this work. However, all else equal, i.e., given two core-selecting payment rules with similar efficiency and fairness properties, we would prefer a rule with higher rather than lower revenue. Still, in our evaluations, a comparison by revenue will generally be secondary.

5. A General Framework for Designing Core-Selecting Payment Rules

In this section we now develop a general framework for designing core-selecting payment rules, that encompasses all of the 84 core-selecting CA payment rules we design in this paper, and subsequently evaluate in a more expansive domain.

5.1. The pricing problem

The pricing problem (finding a particular price vector p) of a core-selecting auction can be formulated as the following mathematical program:

$$\arg \min_p f_1(p, p^{*,1}), f_2(p, p^{*,2}), \dots, f_k(p, p^{*,k}) \quad (13)$$

$$\text{s.t. } p \in \text{Core} \quad \text{Core Constraints \& IR} \quad (14)$$

The functions $f_1, f_2, \dots, f_k : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ are objective functions which minimize some distance between the price vector p and the reference points $p^{*,1}, p^{*,2}, \dots, p^{*,k}$. These objective functions are evaluated in order, i.e., with f_1 being the primary objective, f_2 the secondary, etc. Because this is a sequence of objectives, k programs must be solved in sequence, where only f_k is included in each program and the constraints on p resulting from all $f_i, i < k$, are included as constraints in program k . Equation 14 ensures we produce *core-selecting* prices. While there can in theory be exponentially many core constraints (one for each coalition of bidders), constraint generation typically limits the required number to only a few. Day and Raghavan (2007) were the first to propose a core constraint generation algorithm which made running core-selecting CAs feasible for real-world sized auctions. Recently, Bünz, Seuken and Lubin (2015) further improved this algorithm leading to an additional speed-up.

The reference points $p^{*,1}, p^{*,2}$, etc., used in the objective functions, can be any arbitrary price vectors within or below the core. Multiple authors have discussed the impact of using different reference points and this remains an important open research question (Erdil and Klemperer, 2010; Day and Cramton, 2012). In our experiments, we primarily focus on using VCG payments as the reference point p^* . We additionally also consider $\vec{0}$, as well as reserve prices set by the seller, both of which are independent of bidders' value reports.

As an illustrative example, consider the Quadratic rule (with the MRC), which requires solving a mathematical program with two objective functions: f_1 first minimizes the L_1 distance to $\vec{0}$, and then f_2 minimizes the L_2 distance to the VCG price vector.

5.2. Characterizing Payment Rules By Distance Measures

In general, the objective functions f_1 to f_k can represent arbitrary distance measures between a reference point p^* and the payment vector p . Here, we consider weighted generalizations of the metrics obtained from the Lebesgue L_ρ norms. Following a common slight abuse of terminology we will often refer to the distance metric implied by the L_ρ norm as the L_ρ metric. The weighted form we use is as follows:

$$f_i(p, p^*) = \left(\sum_{j=1}^n \psi(w_j) \cdot |p_j - p_j^*|^\rho \right)^{\frac{1}{\rho}} \quad (15)$$

This is the standard Lebesgue metric, augmented by a set of weights $\vec{w} \in \mathbb{R}^n$ and a simple function $\psi : \mathbb{R} \rightarrow \mathbb{R}$ which is applied to each weight element-wise.¹¹ If a particular payment rule is not guaranteed to result in unique prices in a core-selecting CA setting, then we break ties by further minimizing the L_2 distance to p^* as our final objective.

¹¹Since $\arg \min_x f(x) = \arg \min_x g(f(x))$ for all strictly monotonic functions g , we can in practice minimize $(f_i)^\rho$ in our price optimization instead of f_i . This is particularly helpful since $(f_i)^\rho$ is a convex function.

5.3. Standard Lebesgue Metrics and the Minimum Revenue Core

We first consider a number of unweighted Lebesgue metrics, i.e., where $\psi(x) = x$, and $\vec{w} = \vec{\mathbf{1}}$ in Equation 15. First, we consider the L_2 distance, which gives rise to the *Quadratic* rule we have already looked at in detail. Next we consider the L_∞ distance, which captures the largest element-wise difference between p and p^* . This metric gives rise to the *Threshold* rule introduced by Parkes, Kalagnanam and Eso (2001). Finally, we consider the L_1 distance, which captures the sum of the absolute differences between payments in p and p^* .

Day and Raghavan (2007) argued for selecting prices in the core that are minimal in revenue, the so-called *Minimum Revenue Core (MRC)*. In our framework, this is equivalent to minimizing the L_1 distance to $\vec{\mathbf{0}}$. Most rules used in practice also include a restriction to the MRC. Consequently, for every payment rule, we have studied both a version with and without MRC (i.e., with and without an initial L_1 minimization).

5.4. Weighted Metrics

In addition to the standard L_ρ metrics, we also investigate the weighted metrics introduced by Parkes, Kalagnanam and Eso (2001) in the combinatorial exchange (CE) domain. In a CE, each rule can be defined not only in terms of a distance metric but also in terms of an equivalent algorithmic process. The algorithmic definitions are not available for core-selecting payment rules, because the core constraints bind in ways that are orthogonal to how the algorithms compute prices. However, we can still obtain useful intuitions about the effects of these weighted metrics by appealing to their algorithmic definitions (as captured in their names), as follows:

1. *The Small rule* is defined as an L_1 measure, but with $\psi(x) = x^{-1}$.¹² It favors players with small weights, trying to set their payments to the reference point (or as close to the reference points as the core constraints will allow), while “large” bidders are charged payments as close as possible to their value-reports. This rule has the virtue of maximally removing the strategic incentive from the smallest players. While it may seem unfair to bias the rule for some players and against others, market power does this already so the rule can be viewed as a way to push back against this built-in bias.
2. *The Large rule* is also based on L_1 , but with $\psi(x) = x$, the identity function. It is the inverse of the *Small* rule, i.e., it attempts to set the prices of “large” bidders to the reference point and of “small” bidders to their value-report.
3. *The Fractional rule* also uses $\psi(x) = x^{-1}$, like the *Small* rule, but it uses an L_2 measure instead of L_1 .

5.5. Weights for Weighted Metrics

The Small, Large, and Fractional rules all use weights to bias payments in favor of certain bidders. It remains to identify which weights to use:

1. *VCG Discount Weighting* uses the VCG discount for weights, i.e. $w_i = \Delta_i$. These were the weights introduced in the original Parkes, Kalagnanam and Eso (2001) work, and were chosen because it is ultimately the discounts that are set via the payment rules. However, our results will show that using weights based on *payments*, *reserve prices*, as well as *imputed values* are more effective.

¹²To avoid division by 0, we Winsorize the weights to their smallest non-zero value.

2. *VCG Payment Weighting* uses the VCG payments (calculated based on bidders' reports) as the weights in the mechanism, i.e. $w_i = p_i^{VCG}$. This is a new approach we introduce in this paper, and we will show that it outperforms weighting by discount.
3. *Reserve Price Weighting* uses a set of reserve prices set by the seller as weights, i.e. $w_i = p_i^{RP}$. This approach has been proposed by Baranov (2010), and used in practice in the recent Canadian bandwidth auction (Licensing Framework, n.d.). This can be an effective weighting method, but it heavily relies on the seller actually having reasonable reserve prices, which is often not the case because no good distributional information about bidders' values is available.
4. *Marginal Economy (ME) Weighting* is a novel set of weights we propose in this paper, denoted $w_i = p_i^{ME}$. The idea of marginal economy weighting is to (a) obtain much of the benefit of reserve price weighting, without (b) actually requiring the seller to have good distributional information to compute reserve prices. We now describe this approach in more detail.

The main idea of marginal economy weighting is to compute weights based on an estimate of bidders' values, instead of using bidders' discounts or VCG payments. Because using a bidder's own bid as information about his value for a bundle would open up too many avenues for strategic manipulation, we will instead *impute* an estimate of a bidder's value for his winning bundle, based on the other bidders' bids. The idea is similar to the construction of VCG *payoffs*, which, as we have discussed before, are always equal to a bidder's marginal contribution to the economy. Here, we instead find a reasonable estimate for the *value of the bidder's winning bundle* based on the bids of all other bidders.

Towards this end, we propose the following heuristic: we use the highest value of the bundle in the *marginal economy*, i.e., considering all bidders except the bidder being priced. Such a value has the virtue of not being directly manipulable by the bidder being priced. When the marginal economy contains at least one bid for the same allocation as the bundle being priced, then we can use this losing bid for the marginal economy value. However, this will often not be the case, forcing us to impute a value for the bundle from the bids that *are* in the marginal economy. In the present work, we adopt the following simple heuristic to compute these *imputed values* for *Marginal Economy Weighting*:

$$p_i^{ME}(X_i^*) = \sum_{g \in X_i^*} \max_{j \in N, S \subseteq M: j \neq i, g \in S} \frac{\hat{v}_j(S)}{|S|} \quad (16)$$

That is, for each good g in bidder i 's winning bundle (that needs to be priced), we calculate the maximum per-good value across all marginal-economy bidders. Then we set the weight to be the sum of each of these per-good prices. We find that this weighting function can be highly effective in increasing both incentives and fairness, which we will show empirically in Section 7.

5.6. Parameterized Framework for all Payment Rules

The degrees of freedom outlined in the preceding sections have provided us with 3 weighted metrics (Small, Large, Fractional), 4 weights for these weighted metrics ($\Delta, p^{VCG}, p^{RP}, p^{ME}$), 2 additional non-weighted rules (Quadratic and Threshold), 3 reference points ($p^{VCG}, p^{RP}, \vec{0}$), and the option to use MRC or not for any of our rules. In total, this leads to 84 different core-selecting payment rules:

$$\left(\underbrace{\text{Weighted Metrics}}_3 \cdot \underbrace{\text{Weights}}_4 + \underbrace{\text{(Quadratic, Threshold)}}_2 \right) \cdot \underbrace{\text{Reference Points}}_3 \cdot \underbrace{\text{(MRC, } \overline{\text{MRC}})}_2 = 84$$

Payment Rule Name	Metric	$\psi(w)$	Reference	Metric	$\psi(w)$	Reference	Metric	$\psi(w)$	Reference
Threshold Quadratic w -Large w -Small w -Fractional				f_1			f_2		
		L_∞	$\vec{\mathbf{1}}$	p^*	L_2^\dagger	$\vec{\mathbf{1}}$	p^*		
		L_2	$\vec{\mathbf{1}}$	p^*					
		L_1	w	p^*	L_2^\dagger	$\vec{\mathbf{1}}$	p^*		
		L_1	w^{-1}	p^*	L_2^\dagger	$\vec{\mathbf{1}}$	p^*		
	L_2	w^{-1}	p^*						
MRC Threshold MRC Quadratic MRC w -Large MRC w -Small MRC w -Fractional	f_1			f_2			f_3		
	L_1	$\vec{\mathbf{1}}$	$\vec{\mathbf{0}}$	L_∞	$\vec{\mathbf{1}}$	p^*	L_2^\dagger	$\vec{\mathbf{1}}$	p^*
	L_1	$\vec{\mathbf{1}}$	$\vec{\mathbf{0}}$	L_2	$\vec{\mathbf{1}}$	p^*			
	L_1	$\vec{\mathbf{1}}$	$\vec{\mathbf{0}}$	L_1	w	p^*	L_2^\dagger	$\vec{\mathbf{1}}$	p^*
	L_1	$\vec{\mathbf{1}}$	$\vec{\mathbf{0}}$	L_1	w^{-1}	p^*	L_2^\dagger	$\vec{\mathbf{1}}$	p^*
L_1	$\vec{\mathbf{1}}$	$\vec{\mathbf{0}}$	L_2	w^{-1}	p^*				

Table 2: Here we provide the sequence of metrics and weights that define each of our payment rules. For each rule we use three different reference points $p^* \in \{p^{VCG}, p^{RP}, \vec{\mathbf{0}}\}$. For the weighted rules (Large, Small, Fractional) we use four different weights $w \in \{\Delta, p^{VCG}, p^{RP}, p^{ME}\}$. Note that for the Threshold, Large and Small rules, the final L_2 metric is included only for tie-breaking purposes (as indicated by \dagger).

Table 2 provides an overview of all payment rules we have studied. For each rule, we provide its name, and describe the sequence of metrics and weights that define it.

Remark 2. *Out of the 84 rules described in Table 2, only 9 have been explicitly proposed and studied by other researchers before. We now briefly discuss each of those in turn to put the origins of these rules and the underlying ideas into context.*

Parkes, Kalagnanam and Eso (2001) were the first to suggest using distance-to-VCG-based payment rules in combinatorial settings. They proposed the (1) Threshold rule (without MRC), originally for the CE domain, and later for the CA domain (Parkes, 2002). Then Day and Raghavan (2007) suggested to first enforce prices to be in the MRC, and proposed the (2) MRC Threshold rule. Because minimizing the L_∞ distance to VCG does not guarantee unique prices in CAs, Day and Cramton (2012) refined this approach and proposed the (3) MRC Quadratic rule, which is the rule currently used in practice.

In the original work by Parkes, Kalagnanam and Eso (2001), they also proposed the three weighted rules, but only in combination with the VCG-discount weights, giving rise to the (4) Δ -Large, (5) Δ -Small, and (6) Δ -Fractional rules. Recently, Baranov (2010) proposed to use the Fractional rule in combination with reserve price weighting, giving rise to the (7) p^{RP} -Fractional rule. Finally, Erdil and Klemperer (2010) and Day and Cramton (2012) have both investigated using reference points (in combination with MRC Quadratic) that are independent of all bidders' bids, which includes our reference points p^{RP} and $\vec{\mathbf{0}}$, giving rise to (8) MRC Quadratic($p^ = p^{RP}$) and (9) MRC Quadratic($p^* = \vec{\mathbf{0}}$).*

Subtracting those 9 rules from the 84 rules described in Table 2 leads to 75 new rules which we are the first to study in this paper. Some of these new rules we have designed by simply combining various ideas from prior work (combinations of metrics, weights, reference points, and MRC). However, 36 rules are based on the new weights we have introduced, p^{VCG} and p^{ME} , and we will show in Section 7 that the rules based on those weights perform particularly well. Note that our framework naturally lends itself to extensions via new metrics, weights, reference points, and arbitrary combinations thereof.

Domain	Theoretical BNE	Computational BNE	Std. Dev.
Four-Bidder First-Price Auction	$\frac{n-1}{n} = 0.75$	0.7475	0.0284
Additive Local-Local Global	$\sqrt{8} - 3 \approx -0.1715$	-0.1735	0.0157
Multiplicative Local-Local Global	$\frac{2}{3} \approx 0.6666$	0.6769	0.0177

Table 3: Here we compare theoretical BNE results against the BNE strategies obtained from our algorithm. These include the first-price single item auction (Krishna, 2002), the uncorrelated LLG setting (with additive strategies) (Goeree and Lien, 2014), and the correlated LLG setting (with multiplicative strategies) (Ausubel and Baranov, 2013). Reported values are averages across 50 runs.

6. Methods

6.1. Approximate Bayes-Nash Equilibrium Calculation

To evaluate our rules in realistic settings where closed-form solutions for the bidders’ BNE strategies do not exist or are intractable to derive analytically, we need a computational method for approximating the BNEs. Accordingly, we have developed an algorithm based on the approach by Lubin and Parkes (2009) and earlier work by Vorobeychik and Wellman (2008). Pseudocode of the algorithm is provided in Appendix A.

Our algorithm is based on the general idea of *Fictitious Play*. The algorithm finds outcomes to games by iteratively computing the *expected best response* for a bidder given a set of current strategies for the other bidders. We simplify the problem by defining a fixed number of *bins*, three in our experiments, and assigning each bidder to a bin (*small*, *medium*, or *large*) depending on his value function. We then find a single strategy for each bin that is in *expectation* a best response to the other bidders play. After each iteration, we proceed forward with an affine combination of the best response and the previous strategy. In a BNE, no player (or bin) can improve his utility (in expectation) by playing his best response instead of the BNE strategy. Our algorithm converges if this criterion is met within an error of ϵ . Such an outcome is called a (multiplicative) ϵ -BNE, i.e., each player diverting from the ϵ -BNE can, in expectation, improve his utility by at most a factor of ϵ . Additionally, we ensure that the strategies in the approximated BNE are at most a constant factor δ away from the respective best response. We thus get an (ϵ, δ) -BNE.

More formally, we let $br_i(s_{-i}^*)$ denote agent i ’s best-response to the strategy profile s_{-i}^* of all other agents. To measure the distance between two strategies s_i and s'_i , we use the L_2 -norm $\|s_i - s'_i\|_2$.

Definition 13 (Approximate Bayes-Nash Equilibrium). *A strategy profile $s^* = (s_1^*, \dots, s_n^*)$ is an (ϵ, δ) -Bayes-Nash equilibrium in a sealed-bid auction if, for all bidders i , and all values functions v_i*

$$\mathbb{E}_{v_{-i}} [u_i(s_i^*(v_i), s_{-i}^*(v_{-i}))] \cdot (1 + \epsilon) \geq \mathbb{E}_{v_{-i}} [u_i(br_i(s_{-i}^*), s_{-i}^*(v_{-i}))] \quad \text{and} \quad (17)$$

$$\|s_i^* - br_i(s_{-i}^*)\|_2 \leq \delta, \quad (18)$$

where the expectation in (17) is taken with respect to the distribution over the other bidders’ value functions.

To validate the correctness of our approximate-BNE calculation algorithm, we have computed approximate BNEs for the settings presented in Section 3.5.2 where closed-form expressions are known from the literature. The results are shown in Table 3, which are averages over 50 randomly

initialized runs. We see that our algorithm is able to find approximate BNEs (with additive as well as multiplicative strategies) that are very close to the true BNEs.

To further illustrate the use of our algorithm, we show in Figure 2 the results of an approximate BNE computation with multiplicative strategies for the Quadratic rule, for a more complicated, combinatorial setting with multi-minded bidders, where the bidders are grouped into three bins. Shown are the number of iterations on the x-axis, and the shade factors for each of the three bins on the y-axis. The calculation converges after 23 iterations and finds an approximate BNE with the *Large* bidders shading only a little (0.93), the *Medium* bidders shading a bit more (0.89), and the *Small* bidders shading very much (0.71).




Figure 2: Computation of an approximate Bayes-Nash equilibrium using multiplicative strategies for the Quadratic rule in the Tri-Modal domain.

6.2. Choosing a Strategy Space for our Experiments

In the experiments we present in the next section, we only consider large settings for which no analytical BNE results are known (e.g., 10 bidders, 25 goods, combinatorial preferences, multi-minded bidders). Naturally, the optimal BNE strategies in these settings will be highly complex. To make the approximate BNE calculation computationally feasible we need to choose a reasonably restricted strategy space for our algorithm. For this paper, we limit ourselves to *single-dimensional* strategies to keep the computational run-time manageable (even with single-dimensional strategies, finding the approximate BNE for a single payment rule can take up to one day on a high-performance machine). We study both, (1) *multiplicative* as well as (2) *additive* shading strategies. With multiplicative (shading) strategies, a bidder’s strategy is a shading factor in $[0, 1]$ that he applies multiplicatively to all value reports on all bundles (i.e., shading all values down by the same factor). With additive shading strategies, a bidder’s strategy is an additive factor in $[0, -\infty]$ that he adds to all value reports on all bundles (where adding a negative number also corresponds to shading down his values).

Overall, we obtain very consistent results (in terms of efficiency, fairness, incentives, and revenue) whether we study multiplicative or additive strategies. Thus, for the comparison of our payment rules we could look at the results for *either* multiplicative *or* additive strategies. However, because multiplicative strategies are somewhat easier to interpret (the strategies are always in $[0, 1]$ and can be interpreted independent of the agents’ values), we will discuss our findings by presenting the results for multiplicative strategies in the main body of the paper. The most important results for additive strategies are provided in Appendix F.

We note that in core-selecting auctions, there are reasons to analyze more complex strategies than those we use in our experiments. For example, [Beck and Ott \(2013\)](#) show that *over-bidding strategies* can sometimes be useful. Concretely, they show that by leveraging the core constraints, bidders may be able to decrease their own payment on bundles they expect to win, by bidding more than their true value on bundles they expect to lose. Somewhat orthogonally, [Janssen and Karamychev \(2013\)](#) show that bidders who want to drive up competitors’ payments can use so-called “*spiteful*” *bidding strategies* in combinatorial auctions. This could help a bidder by causing the bidder’s competition to exhaust their budget, or by making the competition overpay for spectrum, hurting the public perception of their position and thus, e.g., their stock price. While we believe that such strategies would be very interesting to study, they are necessarily multi-dimensional, making the equilibrium finding problem significantly more computationally difficult. Thus, as multi-dimensional strategies are outside the scope of our already highly optimized and extremely complex equilibrium solver, we defer the analysis of such strategies to future work.

6.3. Domain Generators

For our experiments, we use two generators to create CA instances from which we evaluate the effectiveness of the payment rules. The first generator we use is the Regions distribution from the Combinatorial Auction Test Suite (CATS) ([Leyton-Brown and Shoham, 2006](#)). CATS has many distributions, but most of them are not particularly suitable for evaluating core-selecting payment rules. Because Regions is based on the geographic allocation of assets, it is the closest match to a spectrum auction environment, where core-selecting payment rules have so far played the biggest role in practice. We therefore use Regions, with the default parameters (16 goods, 8 bidders).

However, we found that even when using the CATS Regions distribution, the average VCG-to-core distance of the auction instances was quite small (i.e., VCG is “almost” in the core). To have greater control over the structure and economic properties of the instances we evaluate, we also constructed our own domain generator based on the one in [Lubin and Parkes \(2009\)](#), and presented in Appendix B. Because this generator produces three different classes of bidders, we call it the *Tri-Modal Generator*. We opted for a generator that explicitly produces this structure because this facilitates the economic interpretation of the results, in particular the *bins* (small, medium, large) to which we assign our bidders in our computational BNE algorithm.

Our generator is designed to be simple, but to still exhibit the key features of combinatorial spectrum auctions. Bidders in the auction are assigned a size, either small, medium or large. Each bidder is only interested in a subset of the goods and constructs a number of XOR bundle-bids by drawing uniform at random from this demand set a number of goods proportional to his assigned size. Bundles are given a value by summing over private per-good values, and then scaling by a factor that is exponential in the size of the bundle above the smallest one demanded by the bidder. We set this exponential to a slightly sub-additive factor, accounting for bidders having a decreasing marginal return for additional goods, once they are able to satisfy their basic business plan.

6.4. Reserve Prices

For a seller to set reasonable reserve prices, she must have good estimates of the bidders’ values. Despite the difficulty of attaining such a distribution in practice, we want to use reserve prices in our experiments, not to increase revenue, as this is not our focus in this paper, but to use them as reference points and weights for our payment rules. Towards this end, we begin by estimating the empirical distribution on bundle values by drawing a large number of bundle-bids from our domain generator. We then take the 20th-percentile of our empirical bundle-value distribution as our bundle-specific

reserve price. To validate this choice, we also evaluated the 10th, 30th and 50th percentile, and found that the 20th-percentile produced both high fairness and revenue, which we therefore adopted for our experiments.

7. Results for Realistically-sized Combinatorial Auction Domains

In this section, we describe the results of our computational BNE analysis, comparing various payment rules according to their (a) efficiency, (b) incentives, (c) fairness, and (d) revenue. Here, we provide results that were produced using our Tri-Modal generator described in Section 6.3. The results are generally consistent with the results for the CATS Regions generator, but they are simply more pronounced for the Tri-Modal generator (which makes reading the results tables easier) than for the Regions generator, because Regions has a small average VCG-to-core distance. The corresponding results for the CATS Regions generator are provided in Appendix E.

Table 4 lists our main results: we compare the various payment rules in our framework when we include the MRC objective, use VCG as the reference point, and adopt multiplicative strategies for the bidders. Results for alternative reference points are provided in Appendix D and additive strategies in Appendix F. In total, we have evaluated each of our 85 rules in eight different ways: four different settings (two domain generators, each with two different VCG-to-Core distances), and with both multiplicative and additive strategies. Computing the BNE for one rule requires approximately 13 hours on a high-performance machine (with 128 GB RAM, two Intel Xeon CPUs at 2.8 GHz, and 10 cores per CPU), with additional 2 hours of time subsequently spent when evaluating the nature of the BNE found by computing the outcome of 100,000 additional game instances in order to produce our efficiency, fairness and revenue results. Across all rules, settings, and strategy spaces, this adds up to more than one year of compute with a single machine, though we have parallelized our computations on a large compute grid.

7.1. Efficiency analysis

The last column of Table 4 is efficiency, from which we can see that all of the rules achieve high efficiency. The lowest efficiency is achieved by the Pay-as-bid rule (98.2%), while all other payment rules lead to an efficiency of 99.4% or higher. Consequently, we immediately turn to our other criteria in evaluating the rules.

7.2. Incentive Analysis

Next, we consider our first main point of interest: the strategies we observe in BNE. We report the strategies for each of the “small”, “medium” and “large” bidders in the second through fourth columns of Table 4, as well as the overall “aggregate incentive” to deviate in the fifth column of Table 4. As was described in Section 4.2, we characterize the latter by the Euclidean distance between truth-reporting and the observed strategies (normalized by value); we color code this column to make the results easier to read.

For the *MRC Quadratic* rule, we observe that the strategies diverge significantly between the “small”, “medium”, and “large” bidders. The small bidders are forced to shade strongly ($\hat{v}_S = .71v_S$) while the medium and large bidders have significantly smaller shades ($\hat{v}_M = .89v_M$ and $\hat{v}_L = .93v_L$). Taken together, these induce a reasonably strong aggregate incentive to deviate at $AI = 31.9$.

Tri-Modal Domain using VCG Reference Point

Payment Rule	Strategies			Aggregate Incentives	% of VCG Payoff			Standard Gini	VCG Gini	% VCG Revenue	Eff.
	Sm.	Med.	Lrg.		Sm.	Med.	Lrg.				
VCG	100	100	100	0.0	100	100	100	43.1	0	100.0	100.0
Pay-as-Bid	28	60	65	89.5	76	108	145	49.1	79.7	92.9	98.2
MRC Threshold	71	88	94	32.4	79	90	105	48.3	61.2	105.9	99.6
MRC Quadratic	71	89	93	31.9	79	90	106	47.5	60.2	105.7	99.7
MRC Δ -Large	65	91	96	36.0	76	87	105	54.0	64.8	107.0	99.4
MRC p^{VCG} -Large	58	88	99	44.1	76	86	113	54.7	70.2	104.9	98.9
MRC p^{RP} -Large	57	86	98	44.8	77	87	118	54.5	70.4	103.5	98.9
MRC p^{ME} -Large	58	86	97	44.3	77	89	118	54.1	70.1	103.3	99.0
MRC Δ -Small	74	84	89	32.8	90	94	108	43.3	59.5	102.1	99.8
MRC p^{VCG} -Small	80	84	87	28.7	90	94	104	42.6	55.1	102.9	99.9
MRC p^{RP} -Small	84	87	88	23.9	91	89	96	40.9	51.3	105.4	100.0
MRC p^{ME} -Small	83	86	89	24.8	91	89	97	42.3	52.2	105.3	100.0
MRC Δ -Fractional	72	87	93	31.8	83	91	105	45.5	59.9	104.9	99.7
MRC p^{VCG} -Fractional	77	86	89	28.8	86	92	104	44.0	56.3	104.2	99.9
MRC p^{RP} -Fractional	82	87	91	23.8	87	86	96	43.5	52.7	106.9	99.9
MRC p^{ME} -Fractional	80	88	92	24.5	85	87	97	44.3	53.5	107.2	99.9

Table 4: The results of a computational BNE analysis with multiplicative strategies (in %). The *Aggregate Incentives*, calculated as $\|s^{\text{truth}} - s^{\text{rule}}\|_2$, indicates the degree of strategizing bidders employ in the BNE of each rule. The *Standard Gini* is calculated over each bidders payoff in equilibrium. The *VCG Gini* is based on the fraction of bidders’ VCG payoffs received in equilibrium.

Among the *Large* rules, this pattern is only made worse, with the small bidders shading even more ($\hat{v}_S \approx .60v_S$) and large bidders being almost truthful. This is expected, given that the *Large* rule aims to provide large bidders with their full VCG discount. This extreme shading by the small players causes the *Large* rules to have overall the poorest aggregate incentives with $AI \approx 40$.

We see that both the *Small* and *Fractional* rules are able to improve the incentives for the small players. In the best case, when using the p^{RP} -*Small* rule, the small players’ strategy is much closer to truth, at $\hat{v}_S = .84v_S$. We note that it is not surprising that even the *Small* rule – which maximally benefits small players – cannot fully compensate for the headwind that small players must endure in combinatorial auctions: the relative distance to VCG is generally higher for small players and the core constraints often have a relatively higher effect on small bidders. Still, the improved incentive effect on small players (while not overly harming the incentives for the other players) yields a corresponding improvement to the aggregate incentives at $AI = 23.9$.

We further observe that the *Fractional* rule also leads to reasonably good incentives for the players compared to the *Quadratic* rule, but not quite as good as the *Small* rule. For example, for p^{ME} -*Fractional*, the incentives to shade vary between $\hat{v}_S = .80v_S$ and $\hat{v}_L = .92v_L$. But importantly, the *maximal incentive* to strategize (here by the small players) as well as the aggregate incentives to strategize are smaller under the p^{ME} -*Fractional* than under the *Quadratic* rule at $AI = 24.5$.

7.3. Fairness analysis

We now turn to our second main point of interest: the comparison of our payment rules in terms of *fairness*. To do this, we begin by computing the fraction of VCG payoff obtained by each winner under the payment rule in each of our 100,000 evaluation instances (shown in columns six through eight of

Table 4).

We next capture two versions of the Gini coefficient for each rule. The first is the “Standard” Gini, calculated on the payoff obtained by each agent in our evaluation instances, as payoff is the natural analogy to “wealth” commonly used when calculating gini coefficients. We present these values in the ninth (and color-coded) column of Table 4. Further, as described in Section 4.3, we believe that the fraction of VCG discount that each participant obtains is also an important reflection of their “fair” apportionment, and so we provide a “VCG”-based Gini in the tenth column (also color-coded). As we shall see, these two measures are remarkably consistent.

The MRC Quadratic rule has a Standard Gini of 47.5 and a VCG Gini of 60.2, which corresponds to its quite unequal distribution of fractions of VCG payoffs: the small players obtain 79% of their VCG payoff while the larger players obtain 106% of their VCG payoff. This is consistent with the motivating example we provided in the introduction, here now verified for a larger setting, and in BNE. Clearly, the MRC Quadratic rule favors the “large” bidders not only on a strategic basis but on a payoff basis as well.

Next we consider the weighted rules, starting with Large. The relative VCG payoffs under Large are even more divergent than under the Quadratic rule and consequently both Gini measures are even higher than under MRC Quadratic. Exactly the opposite occurs, though, under the Small rule, where the unfairness is ameliorated by explicitly favoring the “small” players. Looking at the fractions of VCG payoffs for the Small rule, we observe far greater consistency across bidder sizes than for the Quadratic and Large rules. This translates to the lowest overall Standard and VCG Gini scores of only 40.9 and 51.3 respectively for p^{RP} -Small.

Finally, we turn to the Fractional rule which, when equipped with Δ for weights, seeks to provide the same fraction of VCG payoff across all bidders *ex-post*. While there is no guarantee that targeting optimal fairness *ex-post* will result in a fair rule in equilibrium, we find that Fractional does in fact do well, obtaining a minimum VCG Gini of 52.7 for p^{RP} -Fractional (and a small Standard Gini of 43.5), only slightly higher than that achieved by Small.

Choosing Weights. For each of the weighted rules, we have considered four different weights. The first weight we report is the VCG discount Δ , the weight originally proposed by Parkes, Kalagnanam and Eso (2001). While one might imagine that VCG discount is highly correlated with bidder value, in many domains this correlation is actually low. Indeed, we see in Table 4 that Δ performs surprisingly badly, and in particular is outperformed by all other weights for the Small and Fractional rules. Note that this represents a counterpoint to the results we found for the LLG domain. There, we found that the Δ -Small rule outperformed the Quadratic rule in terms of all dimensions, and in particular in terms of aggregate incentives and fairness. However, here in the larger, more realistically-sized domain, the Δ -Small rule has essentially the same efficiency as the Quadratic rule, slightly better fairness, but somewhat worse aggregate incentives.

Next, we consider our new weights p^{VCG} . We find that using p^{VCG} significantly improves upon using Δ , but is itself dominated, e.g., by reserve prices p^{RP} . In particular, we find that p^{RP} can be extremely effective as weights, for example in the Small rule producing a VCG Gini that is 3.8 lower than that with p^{VCG} weights.

However, in practice, high quality reserve prices will often not be available. Fortunately, our results also demonstrate that for these cases, our new marginal economy p^{ME} weights perform just as well: they produce essentially the same fairness, revenue and efficiency as p^{RP} , but without requiring knowledge of bidders’ value distributions.

This demonstrates how important, not only the distance metric, but also the choice of weights is. In contrast to the Δ -Small rule (which still had worse aggregate incentives than the Quadratic rule), the p^{ME} -Small rule outperforms the Quadratic rule, achieving 100% efficiency, and significantly better fairness (a VCG Gini of 52.2) and aggregate incentives (AI=24.8).

Tri-Model Domain (High-Revenue Quintile) using VCG Reference Point

Payment Rule	Strategies			Aggregate Incentives	% of VCG Payoff			Standard Gini	VCG Gini	% VCG Revenue	Eff.
	Sm.	Med.	Lrg.		Sm.	Med.	Lrg.				
VCG	100	100	100	0.0	100	100	100	34.7	0	100.0	100.0
Pay-as-Bid	25	46	55	103.2	69	87	111	45.6	44.6	115.6	97.5
MRC Threshold	45	77	88	61.1	71	77	90	45.2	39.2	131.2	98.3
MRC Quadratic	45	74	89	62.2	71	78	92	46.3	40.5	129.6	98.0
MRC Δ -Large	38	88	99	63.0	66	63	77	54.2	50.2	145.5	97.1
MRC p^{VCG} -Large	41	81	100	62.1	65	64	86	52.3	48.7	141.5	97.0
MRC p^{RP} -Large	37	74	94	68.4	66	69	98	52.7	49.0	131.2	96.5
MRC p^{ME} -Large	40	76	92	65.6	67	71	97	50.6	46.3	131.7	97.3
MRC Δ -Small	53	62	78	64.5	85	87	97	39.7	33.0	115.6	99.1
MRC p^{VCG} -Small	55	65	80	60.5	83	86	92	40.6	32.5	119.8	99.2
MRC p^{RP} -Small	62	63	79	57.5	86	86	91	39.0	30.9	119.5	99.3
MRC p^{ME} -Small	60	64	81	56.8	86	84	89	39.6	31.7	121.7	99.3
MRC Δ -Fractional	50	68	84	61.4	77	82	94	42.5	35.6	123.9	98.6
MRC p^{VCG} -Fractional	52	67	83	60.5	79	84	92	42.6	34.6	123.2	98.8
MRC p^{RP} -Fractional	60	70	86	52.4	81	79	85	39.9	31.0	128.7	99.2
MRC p^{ME} -Fractional	56	71	87	54.3	78	79	86	41.4	33.0	129.6	99.0

Table 5: The results of a computational BNE analysis with multiplicative strategies (in %). The *Aggregate Incentives*, calculated as $\|s^{\text{truth}} - s^{\text{rule}}\|_2$, indicates the degree of strategizing bidders employ in the BNE of each rule. The *Standard Gini* is calculated over each bidders payoff in equilibrium. The *VCG Gini* is based on the fraction of bidders’ VCG payoffs recieved in equilibrium.

7.4. The VCG-to-Core Distance: Fairness (and Revenue)

Core-selecting pricing rules are often put forward as a way to raise additional revenue in CAs over that available under VCG. While it is true that *ex post* such rules do raise additional revenue, in equilibrium, bidder strategies can often more than compensate for this. As is evident in our results, it is easy for core-selecting rules to produce little or no additional revenue over VCG in equilibrium. For example, in Table 4, the increase in revenue over VCG was on average between 2% and 7%. However, even when the revenue rise is small (or even non-existent), practitioners may still want to use core-selecting rules – not for their revenue properties – but because of the anti-collusion property that the core guarantees. In many real-world settings, providing this guarantee (and avoiding the potential lawsuits that arise from not having it) is more important than the total revenue raised. Ultimately, if high revenue is a concern for the market designer, then techniques such as well-chosen reserve prices should be employed.

This said, we want to be able to evaluate how much revenue the various rules produce in equilibrium. To get a handle on this, we first note that the expected distance between VCG and the MRC *at true values* is a property of the domain, and represents the potential revenue increase of the core over VCG if bidders remained truthful. In equilibrium, the strategic behavior of bidders will reduce the revenue obtained – and the particulars of the payment rule used will mediate how much such reduction occurs. Neither the CATS Regions generator nor our own generator produce instances with a particularly large core-to-VCG distance. Accordingly, we stratify the instances produced by the generators, and consider a domain comprised of only the highest quintile of instances when looking at the truthful distance between the core and VCG.

Table 5 provides results of this type for the Tri-Modal domain. We observe that the overall pattern

of the results is similar to that seen in Table 4. However, because more revenue is available in this setting, comparing the revenue amongst the rules is easier. We can see that Large produces the most revenue, but only through extreme unfairness. In fact, Table 5 shows that fairness and revenue are inverse proportional to each other: the most unfair rule produces the largest amount of revenue, and vice versa. In particular, Small produces significantly less revenue than MRC Quadratic, but is very fair. The Fractional rule, especially in conjunction with marginal economy weights, strikes a balance between fairness and revenue. It produces the same revenue as MRC Quadratic, but with somewhat better efficiency and significantly better fairness.

7.5. Minimum Revenue Core (MRC)

All of the results presented so far are for rules that choose payments from the MRC, as suggested by Day and Raghavan (2007). We have also studied the effect of MRC on the BNE outcome and include a version of Table 4 created without the MRC criterion in Appendix C. Overall, the results with and without MRC are very similar. While relaxing the MRC constraint does change the equilibrium properties slightly for each payment rule, we do not observe any *systematic* shifts in revenue, efficiency, fairness, aggregate incentives, or in the strategies. This seems to suggest that, in Bayes-Nash equilibrium, the effect of the MRC constraint is relatively minor.

7.6. Alternative Reference Points

Additionally, we have also investigated the effect of replacing VCG as the reference point for our rules with our reserve prices, as detailed in Appendix D. When using reserve prices for the reference point, we observe that the differences among our rules narrows. This matches intuition, as the high-quality reserve prices we employ are “doing the work” such that our rules’ weightings have less of an effect. This implies that if good reserve prices are available (which is often not the case), using them may be beneficial. Importantly, however, we observe that the ranking of the various rules remains the same even when using the alternate reference point. Thus, in the case where reserve prices are being employed, the Small and the Fractional rules both remain better choices than the Quadratic rule.

8. Conclusion

In this paper, we have studied the efficiency, fairness, incentive and revenue properties of various payments rules for core-selecting CAs. Our main focus was the *fairness* of these rules in terms of the payoffs of the winners’ in the auction. We have proposed a new fairness concept: a Gini coefficient defined over the fraction of the VCG payoff obtained by the winners. We have then sought core-selecting payment rules that performs well by this fairness measure. To do so, we have defined a framework covering a class of 85 possible rules. We then analyzed these rules using a novel approximate BNE calculation method. We found that all the rules we have evaluated had near perfect efficiency, but the rules varied considerably in terms of their incentive effects, fairness properties and revenue. In particular, we have found that the rule most commonly used in practice, the *Quadratic* rule, strongly favors larger bidders and is thus unfair towards small bidders, while the *Small* rule is the fairest rule.

However, one issue with Small is that it is itself asymmetric between players: small bidders are effectively facing VCG, while the “large” players are effectively playing in a pay-as-bid auction. We find that this asymmetry can help correct for the differential market power that large and small participants have *a priori*. However, in many settings such differential treatment may not be acceptable, depending

on the prevailing policy and law. Fortunately, we have found that the *Fractional* rule provides nearly the same good incentive and fairness properties as the Small rule.

To power weight-based rules like Fractional we have proposed two new weights, including *Marginal Economy* weights, which we find to be almost as effective as reserve prices, but with much lower informational requirements. Moreover, we find that the MRC, typically included in today's core-selecting CAs, has only a minor effect in equilibrium.

Overall, we find that the Quadratic rule is not necessarily the optimal choice for CAs. Instead, considering all of our experimental results, we recommend using the *Marginal-Economy-Fractional* rule, which has the same (or even slightly better) high efficiency and revenue properties as the Quadratic rule, but significantly improves upon it in terms of fairness and aggregate incentives.

In future work, we plan to expand upon the space of strategies that we evaluate in equilibrium, including multi-dimensional (additive and multiplicative) strategies. Additionally, we believe that even more sophisticated heuristics may be developed along the lines of the marginal economy weights we have presented, by using machine learning techniques to interpolate the marginal economy values.

References

- Ausubel, Lawrence, and Oleg Baranov.** 2013. “Core-Selecting Auctions with Incomplete Information.” Working Paper.
- Ausubel, Lawrence, and Paul Milgrom.** 2002. “Ascending auctions with package bidding.” *Frontier of Theoretical Economics*, 1(1): 1–42.
- Ausubel, Lawrence, and Paul Milgrom.** 2006. “The Lovely but Lonely Vickrey Auction.” In *Combinatorial Auctions.*, ed. Peter Cramton, Yoav Shoham and Richard Steinberg, 17–40. MIT Press.
- Baranov, Oleg.** 2010. “Exposure vs. free-riding in auctions with incomplete information.” Working paper.
- Beck, Marissa, and Marion Ott.** 2013. “Incentives for Overbidding in Minimum-Revenue Core-Selecting Auctions.” Working paper, Stanford University.
- Bünz, Benedikt, Sven Seuken, and Benjamin Lubin.** 2015. “A Faster Core Constraint Generation Algorithm for Combinatorial Auctions.” In *Proceedings of the 29th Conference on Artificial Intelligence (AAAI)*. AAAI Press.
- Clarke, Edward.** 1971. “Multipart pricing of public goods.” *Public Choice*, 11(1): 17–33.
- Cramton, Peter.** 2013. “Spectrum auction design.” *Review of Industrial Organization*, 42(2): 161–190.
- Cramton, Peter, Yoav Shoham, and Richard Steinberg.** 2006. *Combinatorial Auctions*. MIT Press.
- Day, Robert, and Paul Milgrom.** 2008. “Core-selecting Package Auctions.” *International Journal of Game Theory*, 36(3): 393–407.
- Day, Robert W., and Peter Cramton.** 2012. “Quadratic Core-Selecting Payment Rules for Combinatorial Auctions.” *Operations Research*, 60(3): 588–603.
- Day, Robert W., and S. Raghavan.** 2007. “Fair Payments for Efficient Allocations in Public Sector Combinatorial Auctions.” *Management Science*, 53(9): 1389–1406.
- Erdil, Aytek, and Paul Klemperer.** 2010. “A New Payment Rule for Core-Selecting Package Auctions.” *Journal of the European Economics Association*, 8(2–3): 537–547.
- Goeree, Jacob, and Yuanchuan Lien.** 2014. “On the Impossibility of Core-Selecting Auctions.” *Theoretical Economics*. Forthcoming.
- Groves, T.** 1973. “Incentives in Teams.” *Econometrica*, 41(4): 617–631.
- Janssen, Maarten, and Vladimir A Karamychev.** 2013. “Gaming in Combinatorial Clock Auctions.” Tinbergen Institute Discussion Paper 13-027/VII.
- Jasso, Guillermina.** 1979. “On Gini’s mean difference and Gini’s index of concentration.” *American Sociological Review*, 867–870.
- Krishna, Vijay.** 2002. *Auction Theory*. Academic Press.

- Kroemer, Christian, Martin Bichler, and Andor Goetzendorf.** 2015. “(Un)expected Bidder Behavior in Spectrum Auctions: About Inconsistent Bidding and its Impact on Efficiency in the Combinatorial Clock Auction.” *INFORMS Group Decision and Negotiation*. Forthcoming.
- Leyton-Brown, Kevin, and Yoav Shoham.** 2006. “A Test Suite for Combinatorial Auctions.” In *Combinatorial Auctions.*, ed. Peter Cramton, Yoav Shoham and Richard Steinberg. MIT Press.
- Licensing Framework. n.d.. *Canadian Government Report*, <http://www.ic.gc.ca/eic/site/smt-gst.nsf/eng/sf10583.html#AnnexE>.
- Lubin, Benjamin, and David Parkes.** 2009. “Quantifying the strategyproofness of mechanisms via metrics on payoff distributions.” In *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence*. UAI.
- Milgrom, Paul.** 2007. “Package Auctions and Exchanges.” *Econometrica*, 75(4): 935–965.
- Othman, Abraham, and Tuomas Sandholm.** 2010. “Envy Quotes and the Iterated Core-Selecting Combinatorial Auction.” In *Proceedings of the Twenty-Fourth Conference on Artificial Intelligence (AAAI)*. Atlanta, GA.
- Parkes, David.** 2002. “On Indirect and Direct Implementations of Core Outcomes in Combinatorial Auctions.” Harvard University.
- Parkes, David C., Jayant Kalagnanam, and Marta Eso.** 2001. “Achieving Budget-Balance with Vickrey-Based Payment Schemes in Exchanges.” In *Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI)*. Seattle, WA.
- Reeves, Daniel, and Michael Wellman.** 2004. “Computing Best-Response Strategies in Infinite Games of Incomplete Information.” In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence (UAI)*.
- Varian, Hal, and Jeffrey MacKie-Mason.** 1994. “Generalized Vickrey Auctions.” University of Michigan.
- Vickrey, William.** 1961. “Counterspeculation, Auctions, and Competitive Sealed Tenders.” *The Journal of Finance*, 16(1): 8–37.
- Vorobeychik, Yevgeniy, and Michael Wellman.** 2008. “Stochastic search methods for Nash equilibrium approximation in simulation-based games.” In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 2*.

APPENDIX

A. The BNE Calculation Algorithm

The fundamental approach we take to solving for the Bayes-Nash Equilibrium is that of *Fictitious Play*. This algorithm finds outcomes to games by iteratively computing the best response for a player given a set of current strategies for the other players. After each such iteration, we update each player’s strategy using a linear combination of the best response and the current current strategy, i.e we proceed forward with a “damped” version of the best response.

Traditionally, such an algorithm would identify a unique strategy for every bidder in the mechanism. However, even with single-dimensional bidder strategies, computing such an equilibrium in our setting is intractable for domains with large numbers of bidders. Consequently, following [Lubin and Parkes \(2009\)](#), we simplify the problem defining a fixed number of *bins*, three in our experiments, and assigning each bidder to a bin. We then find a single strategy per bin, and apply it to all bidders therein. This enables us to compute approximate equilibria, even in games containing many bidders.

For this approach to be effective, “similar” bidders must be assigned to the same bin. While there are many ways to define similarity, we have opted to assign bidders according to their value: our bins are thus defined as the terciles of bidder value (‘small’, ‘medium’ and ‘large’). This would be enough if our bidders were single-minded, i.e. interested in only a single bundle of goods. However, as our bidders are multi-minded, we must define what we mean by their value. We assign them to bins according to the 90th quantile of their value across all of their bundles.

In a BNE no player (or bin) can improve his utility by playing a best-response instead of the BNE strategy. Our algorithm converges if this criterion is met within an error of ϵ . Such an outcome is called an ϵ -BNE, i.e., each player diverting from the ϵ -BNE can, in expectation, gain at most ϵ utility.

For some mechanisms in practice, some players may have utilities that are relatively inelastic with respect to their bids. A sample response utility graph for such a domain is presented in Figure 3. As a consequence, it is possible for the ϵ -BNE condition to be met even while the underlying strategies are not yet stable, because all strategies within the utility plateau are within the support. To account for this, we slightly augment the convergence criterion of our algorithm to require that not only the maximum utility change between the current strategy and the best response is less than ϵ , but also that the maximum change in strategy is less than $\delta \in \mathbb{R}$.

Pseudocode for our approach is provided in Algorithm ???. Our algorithm has a large number of parameters most of which involve a trade-off between validity and runtime. In our experiments we used: $nBidders = 60$, $nSamples = 30$, $nGridPoints = 50$, $nBins = 3$, $\epsilon = .01$, $\delta = 0.015$. To increase the validity of our algorithm we run a final best-response round with increased parameters (e.g. $nBidders = 120$, $nSamples = 60$, $nGridPoints = 100$) and only terminate if the convergence criteria are still met.




Figure 3: Response utilities from a representative step of the BNE calculation from an example of our Tri-Modal generator and the Quadratic rule.

ALGORITHM 1: Iterative Best Response

Data: $nClasses \in \mathbb{N}, \epsilon \in \mathbb{R}^+, \delta \in \mathbb{R}^+$

Result: Class-symmetric ϵ, δ -BNE strategies

// Current strategies

```
1  $S := \vec{\mathbf{1}} \in \mathbb{R}^{nClasses}$ 
  // Best response strategies
2  $\hat{S} := \vec{\mathbf{1}} \in \mathbb{R}^{nClasses}$ 
  // Utilities current strategy
3  $U := \vec{\mathbf{1}} \in \mathbb{R}^{nClasses}$ 
  // Utilities best response
4  $\hat{U} := \vec{\mathbf{1}} \in \mathbb{R}^{nClasses}$ 
5 repeat
6   for  $c = 1$  to  $nClasses$  do
7     // Compute next strategy
8      $S_c = \text{update}(S_c, \hat{S}_c, U_c, \hat{U}_c)$ 
9   end
10  for  $c = 1$  to  $nClasses$  do
11     $\hat{S}_c, \hat{U}_c, U_c = \text{bestResponse}(S, c)$ 
12  end
13 until  $\frac{\hat{U}_c}{U_c} \leq (1 + \epsilon) \wedge |S_c - \hat{S}_c| \leq \delta \quad \forall c \in \{1, \dots, nClasses\}$ 
14 return  $S$ 
```

ALGORITHM 2: Update Rule

Data: $w_{\min} \in \mathbb{R}, w_{\max} \in \mathbb{R}, \alpha \in \mathbb{R}$

Result: Merged strategy

Function $\text{update}(S_c \in \mathbb{R}, \hat{S}_c \in \mathbb{R}, U_c \in \mathbb{R}, \hat{U}_c \in \mathbb{R})$

// Adaptive weight

$w = \arctan(\alpha \cdot (\frac{\hat{U}_c}{U_c} - 1)) \cdot \frac{2}{\pi} \cdot (w_{\max} - w_{\min}) + w_{\min} \in (w_{\min}, w_{\max});$

return $S_c \cdot (1 - w) + \hat{S}_c \cdot w;$

ALGORITHM 3: Best Response Computation

Data: $\delta_{\text{start}}, \delta_{\text{min}}$

Function `bestResponse`($S \in \mathbb{R}^{n_{\text{Classes}}}, c$)

```
// Pattern width
 $\delta = \delta_{\text{start}};$ 
 $S_{\text{center}} = S_c;$ 
while  $\delta > \delta_{\text{min}}$  do
    |  $S_{\text{left}} = S_{\text{center}} - \delta;$ 
    |  $S_{\text{right}} = S_{\text{center}} + \delta;$ 
    |  $S_{\text{best}} = \text{findBest}(S_{\text{left}}, S_{\text{center}}, S_{\text{right}}, c);$ 
    | if  $S_{\text{best}} = S_{\text{center}}$  then
    | | // Narrow pattern
    | |  $\delta = \frac{\delta}{2};$ 
    | end
    |  $S_{\text{center}} = S_{\text{best}};$ 
end
 $U_c = \text{sampleUtility}(S_c);$ 
 $\hat{U}_c = \text{sampleUtility}(S_{\text{center}});$ 
return  $S_{\text{center}}, \hat{U}_c, U_c;$ 
```

ALGORITHM 4: Strategy Comparison

Data: $n_{\min} \in \mathbb{N}, n_{\max} \in \mathbb{N}, c \in \{1, nClasses\}, S \in \mathbb{R}^{nClasses}, v \in (\mathcal{P}(M) \rightarrow \mathbb{R})^{n_{\max}}, A \in \text{Games}^{n_{\max}}$

Function `findBest`($S_{left} \in \mathbb{R}, S_{center} \in \mathbb{R}, S_{right} \in \mathbb{R}$)

```
  n = n_min;
  repeat
    u_left = sampleUtility(S_left, n) ∈ ℝ^n;
    u_center = sampleUtility(S_center, n) ∈ ℝ^n;
    u_right = sampleUtility(S_right, n) ∈ ℝ^n;
    best = arg max_{pos ∈ {left, center, right}} ∑_{i=1}^n u_pos;
    significantDiff=true;
    for other ∈ {left, center, right} \ best do
      // Samples from utility(v, A, S_best, S) – utility(v, A, S_other, S) with v, A i.i.d.
      u_diff = u_best – u_other;
      // Statistical test whether E_{v,A}[utility(v, A, S_best, S) – utility(v, A, S_other, S)] is
      // greater 0 with significance p
      testResult = test(u_diff, n, p);
      significantDiff=significantDiff ∧ testResult
    end
    n = n · 2;
  until significantDiff ∨ n > n_max;
  return S_best;
```

// Produces samples from utility(v, A, S*, S) with v, A i.i.d.

Function `sampleUtility`($S^* \in \mathbb{R}, n \in \mathbb{N}$)

```
  u ∈ ℝ^n = 0;
  for i = 1 to n do
    v_i // random value function drawn from class c
    A_i // random game containing bidder with value function v_i
    // Utility gained under value function v_i when playing S_c in game A_i and other
    // bidders play according to S
    u_i=utility(v_i, A_i, S_c, S);
  end
  return u;
```

Function `test`($u \in \mathbb{R}^n, n \in (n_{\min}, n_{\max}), p \in (0, 1)$)

```
  // Sample mean
  μ = ∑_{i=1}^n u_i / n;
  // Sample Variance
  .5^2 = ∑_{i=1}^n (u_i – μ)^2 / (n – 1);
  // Standard error
  SE = s / √n;
  // Φ^{-1} is the quantile function of the standard normal distribution
  p-hat = Φ^{-1}(-μ / SE);
  return p-hat ≤ p;
```

B. The Tri-Modal Generator Algorithm

Pseudocode for our Tri-Modal generator is provided in Algorithm 5. The generator is careful to capture the property that bidders typically demand only a subset of all goods, producing XOR bundle bids only for sets of goods drawn from this “demandment” set (i.e., demandment : buyer :: endowment : seller). Value for these bundles is calculated as a linear combination of individual good prices, but scaled by a quantity that is exponential in the size of the bundle over the smallest bundle the bidder bids upon. This enables us to model decreasing returns for additional goods, once an bidder’s basic business plan has been met.

For our experiments, we setup the parameters of the generator as follows: $m = 25, n = 9, P_M = .3, P_H = .2, dSizes = [3, 9, 13], nBundles = [2, 4, 8], bSizeL = [1, 4, 7], bSizeH = [2, 6, 9], \gamma = .97$.

ALGORITHM 5: The Tri-Modal Generator

Data: $m, n, P_M, P_H, dSizes, nBundles, bSizeL, bSizeH, \gamma$

Result: n bidder bids, each consisting of XOR bundles

for each bidder do

$size := L$ with $1 - P_M - P_H$, M with P_M and H with P_H

$demandment :=$ draw $dSizes[size]$ goods uniform

 Draw private values $U[0, 1]$, for each $g \in demandment$

for i in $nBundles[size]$ **do**

 Draw $bSize[i]$ from $U[bSizeL[size], bSizeH[size]]$

$bGoods[i] :=$ Draw $bSize[i]$ goods from $demandment$

$bValLin[i] := \sum_{g \in bGoods[i]} private[g]$

end

$sSize := \min_i bSize[i]$

for i in $nBundles[size]$ **do**

$bValNonLin[i] := bValLin \cdot \gamma^{bSize[i] - sSize}$

 Assign bidder $bGoods[i]$ at value $bValNonLin[i]$ as XOR bundle

end

end

C. Results Without MRC

Day and Raghavan (2007) argued for selecting payments that are in the minimum-revenue-core (MRC). In this section we investigate the effect of removing this constraint and select prices that are, perhaps, not in the MRC. We see that the results are very similar to the results with MRC. The MRC constraint does not seem to have a significant effect on the bidder’s strategies, nor on fairness or revenue.

Tri-Modal Domain using VCG Reference Point without MRC

Payment Rule	Strategies			Aggregate Incentives	% of VCG Payoff			Standard Gini	VCG Gini	% VCG Revenue	Eff.
	Sm.	Med.	Lrg.		Sm.	Med.	Lrg.				
VCG	100	100	100	0.0	100	100	100	43.1	0	100.0	100.0
Pay-as-Bid	28	60	65	89.5	76	107	145	49.1	79.7	93.1	98.2
Threshold	68	87	93	35.3	79	90	107	47.9	62.5	105.4	99.5
Quadratic	70	90	94	32.3	78	88	103	47.8	60.3	106.9	99.6
Δ -Large	67	91	97	33.9	73	84	103	55.0	64.4	108.6	99.4
p^{VCG} -Large	54	89	99	47.4	74	83	115	56.1	71.9	105.5	98.7
p^{RP} -Large	51	83	99	52.3	76	90	126	56.3	73.8	100.9	98.4
p^{ME} -Large	52	86	97	50.0	76	89	122	55.3	72.4	102.1	98.7
Δ -Small	70	81	87	38.0	92	97	115	42.5	62.0	99.4	99.7
p^{VCG} -Small	78	81	86	31.9	93	95	107	41.8	57.4	101.5	99.9
p^{RP} -Small	88	84	86	23.8	94	85	94	39.9	53.5	105.8	100.0
p^{ME} -Small	85	84	87	25.6	94	88	97	40.5	53.8	104.8	100.0
Δ -Fractional	70	85	90	34.7	84	93	109	44.1	60.7	103.4	99.7
p^{VCG} -Fractional	77	84	88	30.5	88	93	105	43.2	57.3	103.6	99.9
p^{RP} -Fractional	81	84	88	27.2	90	89	101	40.9	54.0	104.8	99.9
p^{ME} -Fractional	81	87	91	25.1	86	86	96	43.2	53.3	107.1	99.9

Table 6: The results of a computational BNE analysis with multiplicative strategies (in %). The *Aggregate Incentives*, calculated as $\|s^{\text{truth}} - s^{\text{rule}}\|_2$, indicates the degree of strategizing bidders employ in the BNE of each rule. The *Standard Gini* is calculated over each bidders payoff in equilibrium. The *VCG Gini* is based on the fraction of bidders’ VCG payoffs received in equilibrium.

Tri-Model Domain (High-Revenue Quantile) using VCG Reference Point without MRC

Payment Rule	Strategies			Aggregate Incentives	% of VCG Payoff			Standard Gini	VCG Gini	% VCG Revenue	Eff.
	Sm.	Med.	Lrg.		Sm.	Med.	Lrg.				
VCG	100	100	100	0.0	100	100	100	34.7	0	100.0	100.0
Pay-as-Bid	25	46	55	103.2	69	87	111	45.6	44.6	115.6	97.5
Threshold	45	73	89	62.0	70	77	91	46.3	40.2	130.8	98.0
Quadratic	45	75	88	61.7	70	77	90	45.8	39.6	131.1	98.1
Δ -Large	36	89	100	64.7	64	62	77	55.6	51.5	146.5	96.9
p^{VCG} -Large	35	80	100	67.6	63	62	88	54.5	51.6	140.7	96.3
p^{RP} -Large	38	74	94	67.8	64	69	100	52.9	48.8	131.9	96.6
p^{ME} -Large	38	73	90	68.1	66	73	102	51.3	46.9	128.9	97.1
Δ -Small	51	59	76	67.9	87	87	100	40.2	34.1	112.9	98.9
p^{VCG} -Small	54	62	81	62.7	84	84	92	42.1	34.9	119.7	98.7
p^{RP} -Small	70	59	81	54.0	87	81	87	39.9	32.2	122.8	99.0
p^{ME} -Small	67	61	82	54.0	86	80	86	40.5	33.6	124.0	99.0
Δ -Fractional	48	68	82	63.4	78	82	95	41.7	34.9	123.1	98.7
p^{VCG} -Fractional	54	65	84	59.8	79	82	90	42.8	34.7	124.3	98.7
p^{RP} -Fractional	61	66	86	53.8	81	78	86	40.4	31.8	128.1	99.0
p^{ME} -Fractional	58	69	88	53.0	78	77	84	41.9	33.5	131.1	98.9

Table 7: The results of a computational BNE analysis with multiplicative strategies (in %). The *Aggregate Incentives*, calculated as $\|s^{\text{truth}} - s^{\text{rule}}\|_2$, indicates the degree of strategizing bidders employ in the BNE of each rule. The *Standard Gini* is calculated over each bidders payoff in equilibrium. The *VCG Gini* is based on the fraction of bidders’ VCG payoffs received in equilibrium.

D. Results For Alternative Reference Points

Multiple authors (Erdil and Klemperer, 2010; Day and Cramton, 2012) have investigated non-VCG reference points that are independent of the bidders' bids, with the goal of minimizing incentives for strategizing. Here we are showing the reserve reference point, which performs well, but requires sellers to possess significant distributional information on bidders values.

Tri-Modal Domain using Reserve Price Reference Point with MRC

Payment Rule	Strategies			Aggregate Incentives	% of VCG Payoff			Standard Gini	VCG Gini	% VCG Revenue	Eff.
	Sm.	Med.	Lrg.		Sm.	Med.	Lrg.				
VCG	100	100	100	0.0	100	100	100	43.1	0	100.0	100.0
Pay-as-Bid	28	60	65	89.5	76	107	145	49.1	79.7	93.1	98.2
MRC Threshold	76	93	97	25.7	77	80	93	49.4	58.1	110.9	99.7
MRC Quadratic	70	93	98	31.0	78	81	94	50.3	61.3	109.8	99.5
MRC Δ -Large	63	93	98	37.9	76	84	102	55.6	66.2	108.1	99.2
MRC p^{VCG} -Large	62	90	99	39.4	76	82	108	54.5	67.9	107.2	99.1
MRC p^{RP} -Large	63	88	99	38.9	75	83	110	54.5	68.2	106.6	99.1
MRC p^{ME} -Large	61	90	98	40.6	76	83	110	54.4	68.3	106.4	99.1
MRC Δ -Small	78	86	90	28.1	87	89	103	44.5	58.0	104.9	99.9
MRC p^{VCG} -Small	84	86	88	24.5	88	90	100	42.6	53.5	105.1	100.0
MRC p^{RP} -Small	89	88	90	19.3	88	84	92	43.1	51.9	108.1	100.0
MRC p^{ME} -Small	86	88	90	20.8	88	85	92	43.6	51.0	107.7	100.0
MRC Δ -Fractional	78	93	97	23.7	77	79	92	49.4	57.4	111.6	99.7
MRC p^{VCG} -Fractional	82	92	95	20.1	79	79	91	47.2	53.3	111.3	99.9
MRC p^{RP} -Fractional	81	92	96	21.1	79	80	92	47.6	54.4	111.2	99.9
MRC p^{ME} -Fractional	80	91	96	21.8	78	80	92	48.2	55.5	111.1	99.8

Table 8: The results of a computational BNE analysis with multiplicative strategies (in %). The *Aggregate Incentives*, calculated as $\|s^{\text{truth}} - s^{\text{rule}}\|_2$, indicates the degree of strategizing bidders employ in the BNE of each rule. The *Standard Gini* is calculated over each bidders payoff in equilibrium. The *VCG Gini* is based on the fraction of bidders' VCG payoffs received in equilibrium.

Tri-Model Domain (High-Revenue Quantile) using Reserve Price Reference Point with MRC

Payment Rule	Strategies			Aggregate Incentives	% of VCG Payoff			Standard Gini	VCG Gini	% VCG Revenue	Eff.
	Sm.	Med.	Lrg.		Sm.	Med.	Lrg.				
VCG	100	100	100	0.0	100	100	100	34.7	0	100.0	100.0
Pay-as-Bid	25	46	55	103.2	69	87	111	45.6	44.6	115.6	97.5
MRC Threshold	56	82	94	48.0	73	69	77	42.8	37.2	142.7	99.0
MRC Quadratic	54	84	93	49.2	73	69	77	42.9	37.5	142.7	98.9
MRC Δ -Large	38	91	100	62.6	66	61	76	54.8	50.6	147.5	97.1
MRC p^{VCG} -Large	39	84	100	63.4	66	63	84	52.8	49.5	142.6	96.9
MRC p^{RP} -Large	39	80	100	64.8	66	64	87	53.4	50.1	139.7	96.6
MRC p^{ME} -Large	41	80	96	61.9	67	66	89	51.2	46.9	138.6	97.3
MRC Δ -Small	65	67	74	54.9	86	86	96	36.9	27.4	118.8	100.0
MRC p^{VCG} -Small	67	68	80	50.4	84	82	89	38.5	28.9	124.8	99.8
MRC p^{RP} -Small	83	78	59	49.4	84	83	98	39.7	27.4	120.3	99.7
MRC p^{ME} -Small	78	66	80	45.1	83	81	88	38.5	28.6	126.4	99.8
MRC Δ -Fractional	54	86	96	48.3	70	67	75	44.4	39.0	146.3	98.8
MRC p^{VCG} -Fractional	57	84	94	46.7	72	68	77	43.0	36.9	144.1	99.1
MRC p^{RP} -Fractional	56	83	94	47.9	71	69	78	43.3	37.1	142.9	99.0
MRC p^{ME} -Fractional	54	83	95	48.8	71	68	77	43.5	37.9	143.4	98.9

Table 9: The results of a computational BNE analysis with multiplicative strategies (in %). The *Aggregate Incentives*, calculated as $\|s^{\text{truth}} - s^{\text{rule}}\|_2$, indicates the degree of strategizing bidders employ in the BNE of each rule. The *Standard Gini* is calculated over each bidders payoff in equilibrium. The *VCG Gini* is based on the fraction of bidders' VCG payoffs received in equilibrium.

E. Results from the CATS Regions Generator

As noted in Section 6.3, we ran experiments using two generators: the Tri-Modal generator and the CATS Regions generator. In this section, we present the results from the CATS Regions generator. Note that this generator produces auction instances with a significantly smaller VCG-to-core distance, which is why the results are qualitatively similar, but less pronounced than those we have obtained for the Tri-Modal generator.

Regions Domain using VCG Reference Point with MRC

Payment Rule	Strategies			Aggregate Incentives	% of VCG Payoff			Standard Gini	VCG Gini	% VCG Revenue	Eff.
	Sm.	Med.	Lrg.		Sm.	Med.	Lrg.				
VCG	100	100	100	0.0	100	100	100	43.1	0	100.0	100.0
Pay-as-Bid	28	60	65	89.5	76	107	145	49.1	79.7	93.1	98.2
MRC Threshold	71	88	94	32.4	79	89	105	48.3	61.2	105.9	99.6
MRC Quadratic	71	89	93	31.9	79	90	106	47.5	60.2	105.8	99.7
MRC Δ -Large	65	91	96	36.0	76	86	105	54.0	64.8	107.1	99.4
MRC p^{VCG} -Large	58	88	99	44.1	76	85	113	54.7	70.2	104.9	98.9
MRC p^{RP} -Large	57	86	98	44.8	77	87	118	54.5	70.4	103.5	98.9
MRC p^{ME} -Large	58	86	97	44.3	77	88	118	54.1	70.1	103.3	99.0
MRC Δ -Small	74	84	89	32.8	90	93	108	43.3	59.5	102.2	99.8
MRC p^{VCG} -Small	80	84	87	28.7	90	94	104	42.6	55.1	103.0	99.9
MRC p^{RP} -Small	84	87	88	23.9	91	89	96	40.9	51.3	105.5	100.0
MRC p^{ME} -Small	83	86	89	24.8	91	89	97	42.3	52.2	105.4	100.0
MRC Δ -Fractional	72	87	93	31.8	83	90	105	45.5	59.9	104.9	99.7
MRC p^{VCG} -Fractional	77	86	89	28.8	86	92	104	44.0	56.3	104.2	99.9
MRC p^{RP} -Fractional	82	87	91	23.8	87	86	96	43.5	52.7	107.0	99.9
MRC p^{ME} -Fractional	80	88	92	24.5	85	86	97	44.3	53.5	107.2	99.9

Table 10: The results of a computational BNE analysis with multiplicative strategies (in %). The *Aggregate Incentives*, calculated as $\|s^{\text{truth}} - s^{\text{rule}}\|_2$, indicates the degree of strategizing bidders employ in the BNE of each rule. The *Standard Gini* is calculated over each bidders payoff in equilibrium. The *VCG Gini* is based on the fraction of bidders' VCG payoffs received in equilibrium.

Regions Domain using VCG Reference Point without MRC

Payment Rule	Strategies			Aggregate Incentives	% of VCG Payoff			Standard Gini	VCG Gini	% VCG Revenue	Eff.
	Sm.	Med.	Lrg.		Sm.	Med.	Lrg.				
VCG	100	100	100	0.0	100	100	100	43.1	0	100.0	100.0
Pay-as-Bid	28	60	65	89.5	76	107	145	49.1	79.7	93.1	98.2
Threshold	68	87	93	35.3	79	90	107	47.9	62.5	105.4	99.5
Quadratic	70	90	94	32.3	78	88	103	47.8	60.3	106.9	99.6
Δ -Large	67	91	97	33.9	73	84	103	55.0	64.4	108.6	99.4
p^{VCG} -Large	54	89	99	47.4	74	83	115	56.1	71.9	105.5	98.7
p^{RP} -Large	51	83	99	52.3	76	90	126	56.3	73.8	100.9	98.4
p^{ME} -Large	52	86	97	50.0	76	89	122	55.3	72.4	102.1	98.7
Δ -Small	70	81	87	38.0	92	97	115	42.5	62.0	99.4	99.7
p^{VCG} -Small	78	81	86	31.9	93	95	107	41.8	57.4	101.5	99.9
p^{RP} -Small	88	84	86	23.8	94	85	94	39.9	53.5	105.8	100.0
p^{ME} -Small	85	84	87	25.6	94	88	97	40.5	53.8	104.8	100.0
Δ -Fractional	70	85	90	34.7	84	93	109	44.1	60.7	103.4	99.7
p^{VCG} -Fractional	77	84	88	30.5	88	93	105	43.2	57.3	103.6	99.9
p^{RP} -Fractional	81	84	88	27.2	90	89	101	40.9	54.0	104.8	99.9
p^{ME} -Fractional	81	87	91	25.1	86	86	96	43.2	53.3	107.1	99.9

Table 11: The results of a computational BNE analysis with multiplicative strategies (in %). The *Aggregate Incentives*, calculated as $\|s^{\text{truth}} - s^{\text{rule}}\|_2$, indicates the degree of strategizing bidders employ in the BNE of each rule. The *Standard Gini* is calculated over each bidders payoff in equilibrium. The *VCG Gini* is based on the fraction of bidders' VCG payoffs received in equilibrium.

F. Results For Additive Strategies

Here we report our main results for *additive strategies* (while in the body of the paper we have provided the results for multiplicative strategies). Using an additive strategy, a bidder reports his value plus an additive offset. We observe that the results are very consistent with the results we obtained using multiplicative strategies. This consistency (i.e., independency of the choice of strategy space) provides further confidence in our findings regarding the fairness comparison of the various payment rules.

Tri-Modal Domain using VCG Reference Point with MRC and Additive Strategies

Payment Rule	Strategies			Aggregate Incentives	% of VCG Payoff			Standard Gini	VCG Gini	% VCG Revenue	Eff.
	Sm.	Med.	Lrg.		Sm.	Med.	Lrg.				
VCG	0.00	0.00	0.00	0.00	100	100	100	43.1	0	100.0	100.0
Pay-as-Bid	-0.77	-1.13	-1.63	1.20	82	105	138	48.9	80.5	90.8	96.7
MRC Threshold	-0.36	-0.32	-0.28	0.50	83	93	109	47.9	64.5	102.9	99.3
MRC Quadratic	-0.32	-0.31	-0.25	0.46	82	91	106	47.7	62.9	104.3	99.4
MRC Δ -Large	-0.44	-0.28	-0.14	0.60	80	90	109	54.4	69.0	103.5	98.8
MRC p^{VCG} -Large	-0.44	-0.31	-0.04	0.60	78	86	113	53.9	70.6	104.0	98.7
MRC p^{RP} -Large	-0.46	-0.41	-0.11	0.64	80	90	123	53.7	71.6	100.8	98.5
MRC p^{ME} -Large	-0.49	-0.40	-0.17	0.68	80	92	124	53.6	71.9	99.7	98.5
MRC Δ -Small	-0.20	-0.43	-0.47	0.34	90	93	106	43.4	60.5	102.6	99.7
MRC p^{VCG} -Small	-0.20	-0.39	-0.50	0.33	91	93	101	44.5	58.6	103.0	99.8
MRC p^{RP} -Small	-0.14	-0.36	-0.52	0.27	92	89	96	43.4	55.9	104.8	99.8
MRC p^{ME} -Small	-0.17	-0.35	-0.47	0.30	93	89	96	44.1	56.8	104.7	99.8
MRC Δ -Fractional	-0.30	-0.38	-0.40	0.45	87	95	111	44.7	62.7	101.7	99.5
MRC p^{VCG} -Fractional	-0.24	-0.36	-0.43	0.37	88	93	104	45.0	59.2	103.3	99.7
MRC p^{RP} -Fractional	-0.20	-0.35	-0.41	0.32	89	89	99	43.7	56.8	104.9	99.8
MRC p^{ME} -Fractional	-0.22	-0.33	-0.40	0.34	88	90	101	44.3	57.4	104.7	99.7

Table 12: The results of a computational BNE analysis with additive strategies (numbers are additive offsets). The *Aggregate Incentives*, calculated as $\|(s^{\text{truth}} - s^{\text{rule}})/\bar{v}\|_2$, indicates the degree of strategizing bidders employ in the BNE of each rule. The *Standard Gini* is calculated over each bidders payoff in equilibrium. The *VCG Gini* is based on the fraction of bidders' VCG payoffs received in equilibrium.

Tri-Modal Domain (High-Revenue Quantile), VCG Ref. with MRC and Additive Strategies

Payment Rule	Strategies			Aggregate Incentives	% of VCG Payoff			Standard Gini	VCG Gini	% VCG Revenue	Eff.
	Sm.	Med.	Lrg.		Sm.	Med.	Lrg.				
VCG	0.00	0.00	0.00	0.00	100	100	100	34.8	0	100.0	100.0
Pay-as-Bid	-0.85	-1.40	-1.84	1.36	68	75	95	46.1	51.3	123.2	94.8
MRC Threshold	-0.58	-0.60	-0.46	0.83	72	77	89	44.9	41.4	130.5	97.9
MRC Quadratic	-0.58	-0.62	-0.45	0.83	72	77	89	45.2	41.7	130.3	97.8
MRC Δ -Large	-0.66	-0.44	-0.03	0.91	66	66	82	53.4	51.5	140.4	96.7
MRC p^{VCG} -Large	-0.64	-0.59	-0.02	0.89	66	66	89	52.1	50.9	137.0	96.5
MRC p^{RP} -Large	-0.65	-0.69	-0.30	0.92	67	70	101	50.7	48.6	129.4	96.7
MRC p^{ME} -Large	-0.66	-0.66	-0.35	0.94	68	72	99	50.4	48.5	129.0	96.8
MRC Δ -Small	-0.41	-0.99	-1.30	0.76	87	91	102	39.4	33.3	109.7	98.9
MRC p^{VCG} -Small	-0.42	-0.91	-1.21	0.75	86	90	98	39.9	32.4	113.0	99.0
MRC p^{RP} -Small	-0.28	-0.86	-1.43	0.64	90	89	95	37.4	27.8	113.7	99.4
MRC p^{ME} -Small	-0.29	-0.96	-1.06	0.62	88	86	91	38.9	30.7	118.0	99.2
MRC Δ -Fractional	-0.56	-0.80	-0.76	0.85	78	83	96	42.8	38.8	120.5	98.0
MRC p^{VCG} -Fractional	-0.48	-0.73	-0.95	0.76	81	85	93	40.6	33.9	120.6	98.9
MRC p^{RP} -Fractional	-0.43	-0.70	-0.96	0.69	83	83	91	38.1	30.3	122.4	99.2
MRC p^{ME} -Fractional	-0.45	-0.73	-0.59	0.69	79	79	86	41.3	34.9	128.2	98.7

Table 13: The results of a computational BNE analysis with additive strategies (numbers are additive offsets). The *Aggregate Incentives*, calculated as $\|(s^{\text{truth}} - s^{\text{rule}})/\bar{v}\|_2$, indicates the degree of strategizing bidders employ in the BNE of each rule. The *Standard Gini* is calculated over each bidders payoff in equilibrium. The *VCG Gini* is based on the fraction of bidders' VCG payoffs received in equilibrium.

Tri-Modal Domain using VCG Reference Point without MRC and Additive Strategies

Payment Rule	Strategies			Aggregate Incentives	% of VCG Payoff			Standard Gini	VCG Gini	% VCG Revenue	Eff.
	Sm.	Med.	Lrg.		Sm.	Med.	Lrg.				
VCG	0.00	0.00	0.00	0.00	100	100	100	43.1	0	100.0	100.0
Pay-as-Bid	-0.77	-1.13	-1.63	1.20	82	105	138	48.9	80.5	90.8	96.7
Threshold	-0.35	-0.32	-0.28	0.49	82	91	107	47.2	64.2	103.9	99.3
Quadratic	-0.34	-0.31	-0.31	0.48	82	92	108	47.1	63.7	103.7	99.4
Δ -Large	-0.48	-0.27	-0.13	0.65	80	90	110	55.6	70.5	103.1	98.6
p^{VCG} -Large	-0.54	-0.39	-0.06	0.74	78	89	122	55.4	74.1	100.6	98.1
p^{RP} -Large	-0.51	-0.40	-0.06	0.70	78	88	123	55.1	73.2	100.8	98.3
p^{ME} -Large	-0.50	-0.40	-0.14	0.70	79	91	124	54.2	72.8	100.0	98.4
Δ -Small	-0.25	-0.48	-0.53	0.41	93	96	111	43.0	63.5	100.0	99.6
p^{VCG} -Small	-0.19	-0.42	-0.55	0.34	94	92	100	43.9	58.7	102.8	99.8
p^{RP} -Small	-0.11	-0.44	-0.56	0.28	96	87	95	42.7	56.2	104.6	99.8
p^{ME} -Small	-0.12	-0.46	-0.57	0.29	96	89	98	43.5	57.5	103.7	99.8
Δ -Fractional	-0.34	-0.41	-0.44	0.50	88	96	113	44.3	64.7	100.4	99.4
p^{VCG} -Fractional	-0.24	-0.40	-0.45	0.38	90	92	102	44.3	59.7	103.3	99.7
p^{RP} -Fractional	-0.17	-0.41	-0.46	0.31	91	88	98	42.5	56.0	104.9	99.8
p^{ME} -Fractional	-0.20	-0.38	-0.45	0.33	89	90	101	43.1	57.0	104.5	99.8

Table 14: The results of a computational BNE analysis with additive strategies (numbers are additive offsets). The *Aggregate Incentives*, calculated as $\|(s^{\text{truth}} - s^{\text{rule}})/\bar{v}\|_2$, indicates the degree of strategizing bidders employ in the BNE of each rule. The *Standard Gini* is calculated over each bidders payoff in equilibrium. The *VCG Gini* is based on the fraction of bidders' VCG payoffs received in equilibrium.

Tri-Modal Domain (High-Revenue Quantile), VCG Ref. w/o MRC and Additive Strategies

Payment Rule	Strategies			Aggregate Incentives	% of VCG Payoff			Standard Gini	VCG Gini	% VCG Revenue	Eff.
	Sm.	Med.	Lrg.		Sm.	Med.	Lrg.				
VCG	0.00	0.00	0.00	0.00	100	100	100	34.8	0	100.0	100.0
Pay-as-Bid	-0.85	-1.40	-1.84	1.36	68	75	95	46.1	51.3	123.2	94.8
Threshold	-0.58	-0.69	-0.56	0.84	72	78	92	44.7	40.6	127.8	97.8
Quadratic	-0.57	-0.65	-0.48	0.82	71	77	90	44.7	40.7	129.9	97.9
Δ -Large	-0.68	-0.38	0.00	0.93	65	64	81	54.6	52.9	142.8	96.6
p^{VCG} -Large	-0.66	-0.61	-0.02	0.92	64	64	91	52.9	52.1	137.8	96.2
p^{RP} -Large	-0.69	-0.74	-0.28	0.99	66	70	103	52.4	50.6	127.8	96.1
p^{ME} -Large	-0.68	-0.72	-0.37	0.96	67	72	101	51.1	49.3	127.4	96.5
Δ -Small	-0.39	-1.01	-1.45	0.77	88	92	104	39.0	32.8	107.8	98.9
p^{VCG} -Small	-0.45	-0.87	-1.24	0.78	87	89	95	39.6	32.8	114.1	99.0
p^{RP} -Small	-0.23	-1.00	-1.54	0.65	92	88	96	37.9	28.1	112.0	99.1
p^{ME} -Small	-0.24	-1.02	-1.31	0.63	90	86	93	38.3	29.3	114.8	99.2
Δ -Fractional	-0.50	-0.88	-0.88	0.80	79	85	98	41.3	36.1	118.6	98.4
p^{VCG} -Fractional	-0.48	-0.80	-1.12	0.78	82	87	95	40.1	32.9	117.7	98.9
p^{RP} -Fractional	-0.39	-0.86	-0.84	0.67	83	82	89	38.8	30.9	123.4	99.0
p^{ME} -Fractional	-0.42	-0.81	-0.59	0.67	79	79	86	41.4	34.6	128.2	98.7

Table 15: The results of a computational BNE analysis with additive strategies (numbers are additive offsets). The *Aggregate Incentives*, calculated as $\|(s^{\text{truth}} - s^{\text{rule}})/\bar{v}\|_2$, indicates the degree of strategizing bidders employ in the BNE of each rule. The *Standard Gini* is calculated over each bidders payoff in equilibrium. The *VCG Gini* is based on the fraction of bidders' VCG payoffs received in equilibrium. † indicates convergence at the $\epsilon = 2\%$ level without δ .