

2022-12-05

Designing core-selecting payment rules: a computational search approach

B. Buenz, B. Lubin, S. Seuken. 2022. "Designing Core-selecting Payment Rules: A Computational Search Approach" Information Systems Research. <https://doi.org/10.1287/isre.2022.1108>
<https://hdl.handle.net/2144/46237>

"Downloaded from OpenBU. Boston University's institutional repository."



Information Systems Research

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Designing Core-Selecting Payment Rules: A Computational Search Approach

Benedikt Bünz, Benjamin Lubin, Sven Seuken

To cite this article:

Benedikt Bünz, Benjamin Lubin, Sven Seuken (2022) Designing Core-Selecting Payment Rules: A Computational Search Approach. Information Systems Research 33(4):1157-1173. <https://doi.org/10.1287/isre.2022.1108>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2022 The Author(s)

Please scroll down for article—it is on subsequent pages





With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Designing Core-Selecting Payment Rules: A Computational Search Approach

 Benedikt Bünz,^a Benjamin Lubin,^{b,*} Sven Seuken^{c,d}
^aStanford University, Stanford, California 94305; ^bBoston University, Boston, Massachusetts 02215; ^cUniversity of Zurich, 8050 Zürich, Switzerland; ^dETH AI Center, 8092 Zürich, Switzerland

*Corresponding author

Contact: benedikt@cs.stanford.edu,  <https://orcid.org/0000-0003-2082-4480> (BB); blubin@bu.edu,  <https://orcid.org/0000-0002-2786-0997> (BL); seuken@ifi.uzh.ch,  <https://orcid.org/0000-0001-8525-8120> (SS)

Received: January 11, 2020

Revised: February 28, 2021; July 23, 2021

Accepted: October 10, 2021

Published Online in Articles in Advance:
 December 5, 2022

<https://doi.org/10.1287/isre.2022.1108>
Copyright: © 2022 The Author(s)

Abstract. We study the design of core-selecting payment rules for combinatorial auctions, a challenging setting where no strategyproof rules exist. We show that the rule most commonly used in practice, the *Quadratic* rule, can be improved on in terms of efficiency, incentives, and revenue. We present a new computational search framework for finding good mechanisms, and we apply it toward a search for good core-selecting rules. Within our framework, we use an algorithmic Bayes–Nash equilibrium solver to evaluate 366 rules across 31 settings to identify rules that outperform the Quadratic rule. Our main finding is that our best-performing rules are *large-style* rules—that is, they provide bidders with large values with better incentives than does the Quadratic rule. Finally, we identify two particularly well-performing rules and suggest that they may be considered for practical implementation in place of the Quadratic rule.

History: This paper has been accepted for the *Information Systems Research* Special Section on Market Design and Analytics. Ravi Bapna, Martin Bichler, Bob Day, Wolfgang Ketter, Senior Editors; Sasa Pekec, Associate Editor.

Supplemental Material: The online supplement is available at <https://doi.org/10.1287/isre.2022.1108>.



Open Access Statement: This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. You are free to download this work and share with others, but cannot change in any way or use commercially without permission, and you must attribute this work as “*Information Systems Research*. Copyright © 2022 The Author(s). <https://doi.org/10.1287/isre.2022.1108>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by-nc-nd/4.0/>.”

Funding: Part of this research was supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme [Grant Agreement 805542]. This material is based upon work supported by the National Science Foundation [Grant CMMI-1761163].

Keywords: combinatorial auctions • payment rules • core

1. Introduction

The design and implementation of combinatorial auctions (CAs) are a true success story for market design.¹ CAs have found widespread use in practice for selling and buying resources worth billions of dollars. Noteworthy applications include procurement auctions (Sandholm 2013), treasury auctions (Klemperer 2010), and spectrum auctions (Cramton 2013). The advantage of CAs (in contrast to running multiple single-item auctions) is that bidders can express complex preferences over bundles of items, which avoids the exposure problem and can increase efficiency.

There has been a large literature on the design of bidding languages, clearing algorithms, and activity rules for use in CAs (Cramton et al. 2006). However, finding optimal *payment rules* has remained elusive. In this work, we study *direct* payment rules that apply at the end of an auction: these take as input the bidders’

value reports and compute final payments for the winning bidders. Thus, our analysis applies to one-shot, sealed-bid auctions as well as iterative auctions. The best-known real-world application of such payment rules is the *combinatorial clock auction* (CCA), whose *supplementary round* is a sealed-bid CA (Ausubel et al. 2006). The CCA has found widespread adoption in practice. For example, it has been used to conduct more than 15 spectrum auctions, generating more than \$20 billion in revenue (Ausubel and Baranov 2017), and it has been used to auction off offshore wind rights (Ausubel and Cramton 2011).

1.1. Problems with Using a Vickrey–Clarke–Groves Mechanism in Combinatorial Auctions

Our goal is to identify payment rules for CAs that work well in *practice*. This includes optimizing standard mechanism design criteria such as efficiency, incentives,

and revenue. But it also includes paying attention to institutional details and constraints, which may require ruling out certain undesirable outcomes. At first sight, the Vickrey-Clarke-Groves (VCG) mechanism (Vickrey 1961, Clarke 1971, Groves 1973) may seem like an appealing mechanism because it satisfies efficiency, individual rationality, no deficit, and strategyproofness. Unfortunately, for many practical applications, VCG exhibits multiple, severe problems, including the possibility of low-revenue outcomes, incentives for collusion, and incentives for shill bidding (Ausubel and Milgrom 2006).

The fact that VCG may produce very low (and possibly zero) revenue is particularly problematic. To illustrate this, consider a CA with three bidders and two items, A and B . Assume that bidder 1 has value \$10 for the bundle $\{A, B\}$ but zero value for any single item. Assume that bidder 2 has value \$10 for A , value \$0 for B , and value \$10 for $\{A, B\}$. Assume that bidder 3 has value \$10 for B , value \$0 for A , and value \$10 for $\{A, B\}$. VCG allocates A to bidder 2 and B to bidder 3, charging both \$0. Thus, even though there is high competition for the items, the auctioneer's revenue is \$0. This is unfair, because bidder 1 expressed his willingness to pay \$10 for $\{A, B\}$ but won nothing.² In a government auction, it would be unacceptable if public resources were sold for less money than the maximum a group of bidders were willing to offer (Day and Milgrom 2013).

Although reserve prices could, in principle, be used to increase revenue, they cannot avoid the possibility of *noncompetitive* levels of revenue. For spectrum auctions, Ausubel and Baranov (2020b) have shown that low-revenue outcomes have occurred frequently in the assignment stage of the 2016/2017 Federal Communications Commission Incentive Auction, including three zero-revenue outcomes. More recently, researchers as well as practitioners working on ad auctions have also started to address the low revenue produced by VCG when selling complementary ad space (Niazadeh et al. 2021). Thus, the *possibility* of VCG producing very low revenue is of significant practical concern.³

1.2. Minimum Revenue Core-Selecting Payment Rules

At the heart of VCG's problems is that the VCG outcome may be outside 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 the current winners' total payments. A *core outcome* (with respect to reported values) ensures competitive levels of revenue and mitigates bidders' incentives to submit shill bids.

Multiple researchers have proposed mechanisms that guarantee such core outcomes, so-called *core-*

selecting rules (Ausubel and Milgrom 2002; Milgrom 2007; Day and Milgrom 2008, 2013). Note that these rules only guarantee that the outcome lies in the *revealed core* (i.e., with respect to reported values), and this does not guarantee that the outcome also lies in the *true core* (i.e., with respect to the true values). In fact, Goeree and Lien (2016) have shown that, unless the VCG outcome lies in the true core, no payment rule can implement a (true) core outcome in equilibrium. At first sight, this may call into question the value of designing core-selecting rules. However, as Day and Milgrom (2013) have argued, the auctioneer actually wants the core property to hold at the revealed values, such that he can guarantee that there does not exist a group of bidders who have *offered* to pay more in total than what the current winners are paying. Therefore, it has become standard in the research community to use the term “core-selecting” for rules that select outcomes in the revealed core, and we therefore also adopt this terminology in this paper.

Every core-selecting mechanism must select the social welfare-maximizing allocation with respect to reported values (Day and Raghavan 2007). Thus, the quest for core-selecting mechanisms can be reduced to the design of core-selecting *payment rules*. Within the space of core-selecting payments, Day and Raghavan (2007) proposed to only select payments from the so-called *minimum revenue core* (MRC), which is the facet of the core polytope that yields minimal total bidder payments. Day and Milgrom (2013) have shown that MRC-selecting payment rules are Pareto optimal for the bidders. They have further shown that in a Nash equilibrium, MRC-selecting payment rules minimize incentives to misreport, in the sense that they minimize the sum over all individual bidders' maximum gain from deviating from truthful bidding. As this is a desirable property, we also focus our analysis on *MRC-selecting payment rules*.

1.3. The Quadratic Rule

In general, the MRC contains an infinite number of payment vectors, which leaves lots of room for payment rule design. Day and Cramton (2012) proposed the *Vickrey-nearest rule*, also known as the *quadratic rule* (or *QUADRATIC*). *QUADRATIC* selects the payment vector from the MRC that minimizes the Euclidean distance to the VCG payment vector (computed at reported values).

QUADRATIC is the rule used most often in practice (e.g., in the CCA). Certainly, *QUADRATIC* is attractive from a practical point of view, given that it identifies a unique point in the MRC and that this point can be computed via fast quadratic programming techniques. But theory that supports *QUADRATIC* based on its ability to achieve high efficiency or good incentives is

scarce, and we still have an incomplete understanding of its economic properties.

Only recently has the research community started to grapple with the properties of QUADRATIC. For the simple *local-local-global* (LLG) domain, with two items and three bidders, Goeree and Lien (2016) as well as Ausubel and Baranov (2020a) have derived the Bayes–Nash equilibrium of QUADRATIC. It turns out that, even though the rule minimizes the Euclidean distance to VCG, the equilibrium strategies are far from truthful. Multiple researchers have proposed alternative MRC-selecting payment rules (e.g., Erdil and Klemperer 2010 and Ausubel and Baranov 2020a). But they have not been able to identify a rule that dominates QUADRATIC (see Section 2).

1.4. Overview of Our Approach

In this paper, we develop a *computational search approach* for finding MRC-selecting payment rules that outperform QUADRATIC. In terms of performance objectives, we follow prior work on MRC-selecting rules (e.g., Goeree and Lien 2016 and Ausubel and Baranov 2020a) and aim for rules with good efficiency, revenue, and incentives. To capture incentives, we introduce a new measure and prove that it satisfies five economically important desiderata.

The basic idea of our approach is relatively simple: we construct a framework to parameterize the space of MRC-selecting payment rules, and we then algorithmically search a subspace and identify the best-performing rules within this space. However, the details are quite intricate.

First, to make the space of MRC-selecting rules amenable to a computational search, we must choose a well-suited design framework. To this end, we introduce a *parameterized payment rule* we call FRACTIONAL^* , with three parameters: a *reference point* r , a *weight* w , and an *amplification* a (Section 4.2). $\text{FRACTIONAL}^*(r, w, a)$ minimizes the w -weighted Euclidean distance to the reference point r , where the weights can be amplified or dampened by the amplification a . We consider a multitude of reference points, weights, and amplifications, yielding a total of 366 rules.

Although our framework is rich enough to capture all MRC-selecting rules, we restrict our search by requiring that all rules (1) can be computed quickly and (2) can be applied in any CA domain. These goals have guided our selection of reference points and weights (Section 4.3). The outcome of our search is a fixed rule that can be described via the three parameters r , w , and a .⁴

The second challenge relates to comparing the 366 rules in terms of their efficiency, incentives, and revenue. Because no MRC-selecting rule is strategyproof (Goeree and Lien 2016), we cannot simply evaluate the performance of our rules “at truth” (i.e., assuming

that all bidders bid their true values). Instead, we must compute the *equilibrium* for each rule, enabling us to predict bidder behavior and corresponding auction performance if there was a payment rule switch. Given that many high-stakes CAs are only conducted once, and bidders typically keep their valuations secret, the *Bayes–Nash equilibrium* (BNE) is the appropriate solution concept. As deriving 366 BNEs by hand (which typically involves solving a differential equation) is impractical, we use a recently developed *computational BNE solver* by Bosshard et al. (2020) (Section 4.5). Concretely, this is an algorithm that takes as input a payment rule and produces as output an ε -BNE with a very small ε . Given an ε -BNE for each rule, it is then straightforward to compute the relevant measures (efficiency, incentives, and revenue) and evaluate all rules according to their performance.

Our computational approach enables us to evaluate the performance of hundreds of rules in many different settings. For this evaluation, we use two different domain sizes. First, we perform an extensive analysis in the stylized but well-known LLG domain. To enhance the robustness of our results, instead of just studying standard LLG (with uniform distributions and no correlation), we create 29 different variations of the LLG domain (varying, e.g., the marginal distribution as well as the correlation between bidders); we call each such variation a *setting*. Thus, with 366 rules evaluated on 29 settings, we study 10,614 rule-setting combinations in LLG. This allows us to identify a set of 20 very good *all-rounder* rules—that is, payment rules that outperform QUADRATIC (on average) across all 29 LLG settings in terms of efficiency, incentives, and revenue. Seven of these rules even dominate QUADRATIC in *every* dimension in *every* LLG setting. To see how well rule performance generalizes to a different domain, we then select seven of the best-performing rules from the first step and also evaluate them in the larger and more complex LLLLGG domain. We find that those rules also perform well in LLLLGG.

As a final step in our analysis, we aim to identify patterns to answer the question, what makes a rule a good rule? An interesting pattern we identify is that our best-performing rules are *large-style* rules in that they provide better incentives for bidders with larger values. By contrast, Lubin and Parkes (2009) found *small-style* rules performed well in the budget-balanced combinatorial exchanges (CEs) they studied. This suggests that the structural differences between CEs and CAs may be important in the design of their payment rules.

A last finding we want to emphasize is that there is not one best reference point or one best set of weights. Instead, our results show that the *combination* of the three parameters matters. Our analysis leads to a set

of well-performing rules, where for each rule, the reference point, the weights, and the amplification are perfectly tuned to complement each other. We highlight two rules that stand out for their strong performance across all settings and their simplicity, which makes them a good candidate for use in practice in place of QUADRATIC.

Overall, our paper shows that computational search can be a powerful tool for mechanism design. Furthermore, our results demonstrate that large improvements over QUADRATIC are possible. We hope that some of our best-performing rules will spur new research and that we have also enlarged the space of payment rules that may be considered for implementation in practice.

2. Related Work

Research on the design of mechanisms that achieve some form of “approximate incentive compatibility” has a long history. Parkes et al. (2001) introduced the idea of finding prices that minimize the distance to VCG, first for combinatorial exchanges and later for CAs (Parkes 2002). Building on this, Day and Raghavan (2007) proposed to first minimize the total revenue before minimizing the distance to VCG. These papers can be seen as precursors to the design of the QUADRATIC rule.

Erdil and Klemperer (2010) argued that non-VCG reference points that are independent of the bidders’ reports offer better incentives at the margin of truthful play. However, they do not offer a concrete payment rule, nor do they offer an argument about what happens when further deviations are necessary. Along the same lines, Day and Cramton (2012) studied QUADRATIC with a zero reference point (or ZERO for short) instead of VCG. Using computational experiments (simulating truthful bidding), they found that the use of ZERO favors higher-valued bidders. However, they did not study this rule in a Bayes–Nash equilibrium, nor did they analyze its efficiency.

Ausubel and Baranov (2020a) provided an analytical study of three MRC-selecting payment rules (QUADRATIC, proxy, and nearest-bid), varying the distributional assumptions and the degree of risk aversion. They found that MRC-selecting payment rules perform better in terms of efficiency and revenue when bidders’ values are more correlated, whereas VCG performs worse. However, they did not identify a new payment rule with superior properties to QUADRATIC.

Marszalec (2018) compared three core-selecting payment rules against VCG via laboratory experiments. He finds relatively low efficiency under VCG, which may be explained by attempts to collude by the bidders. For the three core-selecting payment rules, he found that they lead to higher efficiency than VCG,

even when the bidders do not follow the equilibrium bidding strategies.

Parkes et al. (2001) already proposed three *weighted payment rules*. In particular, they suggested using the bidders’ VCG payoffs to influence which core point to select. Ausubel and Baranov (2017) reported that weighted versions of QUADRATIC have been used in recent CCAs conducted in Australia and Canada. However, they did not use VCG payoff but well-chosen reserve prices to power their weighted version of QUADRATIC. The (intuitive) reason in favor of using reserve-price weights is that those reserve prices are not manipulable, whereas the VCG payoff is. However, no theoretical analysis of this reserve-price-weighted version of QUADRATIC exists.

3. Preliminaries

3.1. Formal Model

In a CA, there is a set M of m distinct, indivisible items, and a set N of n bidders. Each bidder i has a *valuation* v_i that, for every bundle of items $S \subseteq M$, defines bidder i ’s value $v_i(S) \in \mathbb{R}_{\geq 0}$ (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 generalizes to sellers with values (see Day and Cramton (2012) for specifying reserve prices).

We let $\mathbf{p} = (p_1, \dots, p_n)$ denote the payment vector, with p_i denoting bidder i ’s payment. We assume that bidders have quasilinear utility functions of the form $u_i(S, p_i) = v_i(S) - p_i$. Bidders make reports about their values for bundles to the mechanism, denoted by $\hat{v}_i(S) \in \mathbb{R}_{\geq 0}$, which may be nontruthful. We let \hat{v} denote the tuple of all bidders’ value reports. We follow the majority of prior work (Goeree and Lien 2016, Ausubel and Baranov 2020a, Bosshard et al. 2020) and assume that bidders only bid on bundles they are directly interested in (i.e., where removing any item from the bundle would strictly decrease the value).⁵

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., $X_i \cap X_j = \emptyset \forall i, j \in N$). In addition to the allocation rule, a mechanism also specifies a *payment rule*, which determines the payment vector $\mathbf{p} = (p_1, \dots, p_n)$. Together, these define the *outcome* $O = \langle X^*, \mathbf{p} \rangle$. An outcome O is called *individually rational* if $\forall i \in N: u_i(X_i^*, p_i) \geq 0$.

3.2. VCG Payments and VCG Payoff

Next, we define two auxiliary concepts based on the well-known VCG mechanism (Vickrey 1961, Clarke

1971, Groves 1973), which we later use as components in the design of our payment rules.

Definition 1 (VCG Payments). Given an allocation X^* and bidders' value reports \hat{v} , bidder i 's VCG payment is defined as $p_{VCG,i} = \sum_{j \neq i} \hat{v}_j(X^{-i}) - \sum_{j \neq i} \hat{v}_j(X^*)$, where X^{-i} is the welfare-maximizing allocation when all bidders except i are present.

Note that VCG payments can be computed even if a different payment rule than VCG is used; in this case, the VCG payments are the payments the bidders *would have paid* if the VCG payment rule *had been used* but applied to the value reports submitted under the actual payment rule. We use those VCG payments in the definition of several of our more sophisticated payment rules.

Analogously, we also define a bidder's (reported) VCG payoff, which is the payoff the bidder *would have gotten*, given his reported value, if the VCG payment rule *had instead been used* in place of the actual payment rule. Note that the payoff may be different from the bidder's utility, which is always evaluated at the *true* values.

Definition 2 (VCG Payoff). Given allocation X^* and bidders' value reports \hat{v} , bidder i 's (reported) VCG payoff is his reported value minus his VCG payment: $\pi_{VCG,i} = \hat{v}(X_i^*) - p_{VCG,i}$.

3.3. Bayes–Nash Equilibrium

We assume each bidder i knows his own valuation v_i but only has distributional information about all other bidders' valuations. We assume that bidders' valuations are drawn from a joint distribution with probability density function (PDF) $f: \mathbb{R}^{n \cdot 2^m} \mapsto \mathbb{R}_{\geq 0}$, and that this distribution is common knowledge. Thus, from each bidder's perspective, the auction is a game of incomplete information which is why BNE is the appropriate solution concept.

We let s_i denote bidder i 's strategy, which is a mapping from his true valuation v_i to a possibly nontruthful report \hat{v}_i . Given a valuation and a strategy from each bidder, this determines the outcome of the auction. We 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 valuations of all bidders except i , and analogously for the strategies s_{-i} . We say that a strategy profile s^* is an ε -BNE if no bidder has a profitable deviation from this strategy profile netting him more than ε in utility.

Definition 3 (ε -Bayes–Nash Equilibrium). A strategy profile $s^* = (s_1^*, \dots, s_n^*)$ is an ε -Bayes–Nash equilibrium (ε -BNE) if, for all bidders $i \in N$, for all valuations v_i :

$$\mathbb{E}_{v_{-i}}[u_i(s_i^*(v_i), s_{-i}^*(v_{-i}))] \geq \mathbb{E}_{v_{-i}}[u_i(\hat{v}_i, s_{-i}^*(v_{-i}))] - \varepsilon$$

for all possible reports \hat{v}_i , (1)

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

In this work, we use numerical algorithms with limited precision to find an equilibrium. We therefore adopt ε -BNEs as our solution concept, where ε is a suitably small constant fixed a priori.

4. A Computational Search for MRC-Selecting Payment Rules

In this section, we present our design framework as well as our computational search approach.

4.1. MRC-Selecting Payment Rules

As discussed in Section 1.2, we follow prior work (e.g., Day and Milgrom 2013) and aim for payments in the revealed core, which we simply call the *core* going forward.

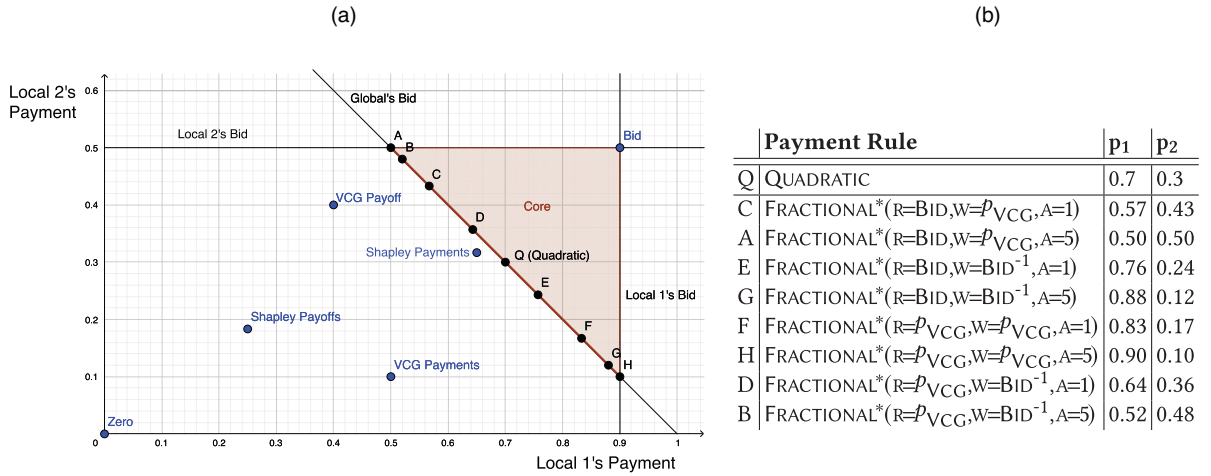
Definition 4 ((Revealed) Core). We let \hat{v} denote the value reports, X^* denote an optimal allocation, W denote the set of winners, $C \subseteq N$ denote a coalition of bidders, and X^C denote an optimal allocation that would be chosen by the mechanism if only the bidders in the coalition C were present. A payment vector p is in the (revealed) core if, in addition to individual rationality, the following set of core constraints hold:

$$\sum_{i \in W \setminus C} p_i \geq \sum_{i \in C} \hat{v}_i(X^C) - \sum_{i \in C} \hat{v}_i(X^*) \quad \forall C \subseteq N. \quad (2)$$

Enforcing payments to be in the core puts lower bounds (i.e., 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 (at the reported values) there exists no coalition that is willing to pay more to the seller than the current winners' payments.

Figure 1(a) illustrates the core in the so-called LLG domain with two items and three bidders. In this example, local bidder 1 bids \$0.9 for item A , local bidder 2 bids \$0.5 for item B , and the global bidder bids \$1 for the bundle $\{A, B\}$. In the figure, the x axis corresponds to local bidder 1's payment, and the y axis corresponds to local bidder 2's payment. The core is the red shaded area. The MRC is the facet of the core polytope that yields the minimal total bidder payments (Day and Raghavan 2007). In Figure 1(a), the MRC is indicated by the bold red diagonal line. A rule that always selects a payment vector in the MRC is called *MRC selecting*.

In Figure 1(a), we also depict several payment vectors of particular importance in blue (including p_{VCG} —i.e., VCG payments), which we motivate and define formally in Section 4.3. Finally, the points A–H and Q are the payments chosen by nine different MRC-selecting payment rules. The point Q is the payment vector chosen by QUADRATIC. For this rule, p_{VCG}

Figure 1. (Color online) Example of the Core and MRC Payments

A graphical depiction of the *core* and the reference points p_{VCG} , π_{VCG} , $p_{SHAPLEY}$, $\pi_{SHAPLEY}$, ZERO, and BID. We also show the payments chosen by 8 different rules as well as QUADRATIC.

Payments for QUADRATIC and 8 other rules constructed from reference points $\{BID, p_{VCG}\}$, weights $\{p_{VCG}, BID^{-1}\}$ and amplifications $\{1, 5\}$.

serves as the *reference point* of the payment rule, which means that the rule selects a point in the MRC that is defined relative to this point. Specifically, QUADRATIC selects the point in the MRC that minimizes the Euclidean distance to the reference point p_{VCG} (Day and Cramton 2012).

4.2. Design Framework for New MRC-Selecting Payment Rules

We now introduce a framework that allows us to concisely define all MRC-selecting payment rules we consider in this paper. The framework requires three parameters: (1) a *reference point function* $r(\cdot)$ that maps bidders' reports \hat{v} to a payment vector, (2) a *weight function* $w(\cdot)$ that maps \hat{v} to a weight vector, and (3) an *amplification scalar* $a \in \mathbb{R}_{\geq 0}$.

Definition 5 (Algorithmic Framework for MRC-Selecting Payment Rules). Given a vector of reported values \hat{v} , a *reference point function* $r(\cdot)$, a *weight function* $w(\cdot)$, and an *amplification scalar* a , the unique MRC-selecting payment vector p^* is chosen to be

1. within the *minimal revenue core* according to \hat{v} , and
2. within this, the payment vector that minimizes the weighted and amplified Euclidean distance to the reference point:

$$p^* = \arg \min_{p \in \text{MRC}} \sqrt{\sum_{i=1}^n \frac{|p_i - r_i(\hat{v})|^2}{w_i(\hat{v})^a}}. \quad (3)$$

We adopt the standard terminology and simply refer to a *reference point* r and *weight* w , leaving the dependence on the bidders' reports implicit. We also refer to r_i as bidder i 's reference point.

Our framework generalizes QUADRATIC: it allows for any reference point r and uses a *weighted* Euclidean distance, where the weights w can be *amplified* by raising them to a fixed power a .⁶ Given that the weights are in the denominator, the minimization in Equation (3) has the effect that the payment of a bidder with smaller weights is closer to his reference point. As $a \rightarrow \infty$, the effect of the weights get amplified. For $a \in [0, 1)$, the weights get dampened, completely removing the effects of the weighting for $a = 0$. The framework naturally encompasses QUADRATIC by choosing p_{VCG} as the reference point, setting the weights to EQUAL = 1, and setting the amplification to 1.

The usefulness of our framework comes from the fact that any MRC-selecting rule can now simply be defined via a triple (r, w, a) , which allows us to search for good rules by searching a three-dimensional parameter space. In general, our framework is rich enough to capture any MRC-selecting rule (by choosing suitable, but arbitrarily complex reference points). However, we limit our search to rules (and thus reference points and weights) that (1) can be computed quickly and (2) can be applied in any CA domain. With the second point we mean that the rule is well defined for any CA instance (with any number of items or bidders) and can therefore immediately be applied in practice.⁷

Our framework generalizes the FRACTIONAL rule introduced by Parkes et al. (2001). We therefore call all of our rules FRACTIONAL*. For example, FRACTIONAL* (R = p_{VCG} , W = π_{VCG} , A = 5) refers to the rule we obtain from Equation (3) with reference point p_{VCG} , weight π_{VCG} , and amplification of 5. Note that we use uppercase "R," "W," and "A" when writing out the full names of the rules to enhance readability.

4.3. Instantiating Our Framework: 366 MRC-Selecting Payment Rules

We now present the parameter sets we consider within our FRACTIONAL* framework when searching for new MRC-selecting rules. Recall that we use p for payments and π for payoffs. We use the -1 superscript to denote taking the *reciprocal* of the weights, which reverses the prioritization they construct. Because we include reference points that are strictly inside the core, we consider *mirroring* all reference points that are inside the core across the nearest MRC facet (denoted by the “M” superscript). This ensures that all reference points are outside and below the core, which implies that the direction to the MRC line (and thus the sign of the expression) is always consistent. By BID , we refer to the bid vector of the winning bidders as charged by the simple first-price payment rule. By p_{SHAPLEY} and π_{SHAPLEY} , we refer to payment and payoff vectors computed based on the well-known Shapley value, respectively (see Online Supplement A for formal definitions).⁸ The following is the list of reference points, weights, and amplifications that we consider in our analysis:

1. *Reference points:* ZERO, BID, BID^M , p_{VCG} , p_{SHAPLEY} , and p_{SHAPLEY}^M
2. *Weights:* EQUAL, BID, BID^{-1} , π_{VCG} , π_{VCG}^{-1} , p_{VCG} , p_{VCG}^{-1} , π_{SHAPLEY} , $\pi_{\text{SHAPLEY}}^{-1}$, p_{SHAPLEY} , and p_{SHAPLEY}^{-1}
3. *Amplification:* 0.5, 1, 2, 3, 5, 10

We choose these specific parameter sets such that we cover all previously established rules (enabling for comparisons against the literature), rules close by, and many new rules.⁹ In particular, our search includes all previously established reference points (ZERO, BID, p_{VCG} , and p_{SHAPLEY}) and weights (EQUAL and π_{VCG}). The mirroring and inversion operations lead to two new reference points and five new weights. Considering Shapley-based weighting is a new idea, which we discuss further in Online Supplement A. Finally, the set of amplification factors are chosen such that we explore not only the effect of dampening the influence of the weights (with $A=0.5$) but also, more importantly, the influence of extremizing the effect of the weights (with $A \in \{2, 3, 5, 10\}$). Taking the cross-product of these parameter sets results in 366 *different MRC-selecting rules*, which stands in stark contrast to the roughly 15 MRC-selecting rules that have been studied in prior work.¹⁰

Even though we do not explore the infinitely large set of *all* possible MRC-selecting rules, our set of rules provides us with good *coverage* of the MRC. To gain some intuition, consider Figure 1, where we see that the payment vectors determined by the nine MRC-selecting rules in this example already provide good coverage of the MRC. In Online Supplement B, we provide a detailed analysis of 10,000 random auction

instances, showing that our full set of 366 rules covers the MRC very densely for almost all auction instances. Given that rules only differ in which points on the MRC they select, this shows that we consider a very large and diverse set of rules.

4.4. Reserve Price-Weighted MRC-Selecting Rule from the 2019 Canadian Auction

Recently, some of the auctions run in practice have started to experiment with *weighted* MRC-selecting rules. Most prominently, the Government of Canada (2019) used a weighted MRC-selecting rule with reference point p_{VCG} to run a spectrum auction, even though there is no prior work studying the properties of this rule. As weights, the rule used *reserve prices*, which were set by the auctioneer. For comparison purposes, we include such a reserve price-weighted rule in our analysis. As in the Canadian auction, we use per-item reserve prices as weights for the rule. The weight used for a bidder is then the sum of the reserve prices of the items won by this bidder. Unlike in the Canadian auction, we do not enforce the reserve prices themselves but rather just use them to compute the weights in the payment rules.¹¹ We do this because none of our other rules use reserve prices, and doing so for one rule would distort the revenue comparison. In LLG, this rule turns out to be equivalent to QUADRATIC because if the local bidders win, they both win exactly one item (and we use the same reserve price for all items). In LLLLGG, however, bidders can win a different number of items such that the reserve price-based weights are distinct and have an effect. Therefore, we include this rule in our analysis in Section 7.

4.5. Searching for Optimal Rules via a Computational BNE Solver

Given our framework for the design of MRC-selecting payment rules, we next consider how to search through the space of *candidate rules* to find one with high efficiency, good incentives, and high revenue *in BNE*. To this end, we employ a BNE-finding algorithm that was recently introduced by Bosshard et al. (2017, 2020). See Online Supplement C for details on the algorithm. We provide the source code for our use of this algorithm at <https://github.com/marketdesignresearch/Comp-Search-BNE>.

As this BNE algorithm is a numerical algorithm with limited precision, it only finds ϵ -BNEs for some $\epsilon > 0$. For LLG, we only report rules for which we find a 0.001-BNE (where the ϵ of 0.001 corresponds to 0.1% of any nontruthful bidder’s maximum value). In LLLLGG, we only report MRC-selecting rules for which we find a 0.01-BNE (where the ϵ of 0.01 corresponds to 0.5% of any bidder’s maximum value). Computing these high-precision ϵ -BNEs takes minutes to hours in LLG and multiple days in LLLLGG. For this

reason, LLLLGG is currently the most complex domain we can feasibly study with this approach.

Equipped with the BNE solver, we perform an exhaustive search over the discretized parameter space. We compute the ε -BNEs of every rule for every setting we consider. To find good rules, we compare all rules according to multiple performance measures, which we detail in Section 5. We provide open-source access to the code used in our experiments at github.com/marketdesignresearch/Comp-Search-BNE.

5. Design Dimensions and Performance Measures

In this section, we introduce formal measures for the three goals we strive for when designing MRC-selecting payment rules: (1) high efficiency, (2) high revenue, and (3) good incentives.

5.1. High Efficiency

From a social planner’s perspective (e.g., a government auctioning off spectrum), it is desirable to maximize the *social welfare* generated by the allocation under a mechanism (i.e., the sum of the winners’ true values for their allocations). The *efficiency* of a mechanism in a given auction instance is defined as the fraction of the social welfare that the mechanism achieves in that instance relative to the maximum possible. When measuring the performance of a mechanism for a distribution of instances, one must generalize this standard concept. To this end, let I denote an auction instance, which in our analysis simply corresponds to a valuation profile v . Recall that the instances are drawn from a known joint probability distribution with PDF f . Let $SW_{\text{OPT}}(I)$ denote the social welfare obtained under the optimal allocation, evaluated at truth. Let $SW_M(I)$ denote the social welfare obtained by the mechanism M when all bidders play their BNE strategies, evaluated at truth. We define the efficiency of mechanism M given distribution f as

$$\text{EFFICIENCY}_{M,f} = \frac{\mathbb{E}_{I \sim f}[SW_M(I)]}{\mathbb{E}_{I \sim f}[SW_{\text{OPT}}(I)]}. \quad (4)$$

In other words, our measure is the expected welfare of the mechanism divided by the expected welfare of the optimal allocation. This is a standard definition for efficiency used in prior work (e.g., by Goeree and Lien (2016)).¹² When computing Equation (4) in our experiments, we use numerical integration (in LLG) or Monte Carlo sampling (in LLLLGG).

5.2. High Revenue

A primary motivation for using MRC-selecting payment rules is that they achieve high enough revenue such that no subset of bidders reported to be willing to pay more than the current winners’ total payments.

Note that all MRC-selecting rules satisfy this. However, all else equal, many auctioneers prefer higher over lower revenue (Ausubel and Baranov 2017). Thus, we also include revenue in our list of desiderata. Let $\text{REVENUE}_{\text{VCG}}(I)$ denote the revenue (i.e., the sum of all winners’ payments) obtained under VCG, evaluated at truth. Let $\text{REVENUE}_M(I)$ denote the revenue of mechanism M when all bidders play their BNE strategies, evaluated at truth. To measure revenue, we again follow Goeree and Lien (2016) and define a measure analogous to the one for efficiency:

$$\text{REVENUE}_{M,f} = \frac{\mathbb{E}_{I \sim f}[\text{REVENUE}_M(I)]}{\mathbb{E}_{I \sim f}[\text{REVENUE}_{\text{VCG}}(I)]}. \quad (5)$$

We use a relative measure (normalizing by the VCG revenue) to enable comparisons across rules and settings. Note that with this measure, a rule can achieve more than 100% revenue.

5.3. Good Incentives

Finally, we seek payment rules with “good incentives.” Even though QUADRATIC is not strategyproof, auction designers have argued in favor of it because it “induces truthful bidding” (Cramton 2013, p. 165) and because it “minimizes the bidders’ ability to benefit from strategic manipulation” (Day and Raghavan 2007). One argument is that, if the rule is “approximately strategyproof,” then finding a beneficial deviation from truthful bidding may be so hard that many bidders may just report truthfully (Day and Milgrom 2008). Of course, because there is no strategyproof MRC-selecting payment rule, there will always remain some strategic opportunities for the participants; however, we would like these opportunities in BNE to be as small as possible.

Several different measures of approximate incentive compatibility have been considered in the literature in different contexts (e.g., Lubin and Parkes 2012, Balcan et al. 2019, Deng and Lahaie 2019, and Balseiro et al. 2021). However, these definitions consider unilateral deviations from a truthful strategy profile. By contrast, we need to evaluate the distance to truthful reporting in BNE, and we thus need a different measure that captures this. To this end, in Online Supplement D.1, we introduce five desiderata that we argue an economically useful incentive measure should satisfy. We then construct a new incentive measure step by step and prove that it satisfies all five desiderata (see Proposition 1 in Online Supplement D.2). We now present the resulting measure.

As for efficiency and revenue, we measure the incentives of a mechanism given a distribution f of bidders’ values. This is necessary because the manipulability of a payment rule depends on a bidder’s value. In particular, for some rules, a bidder may benefit more from manipulating if he has a small value,

whereas for other rules, a bidder may benefit more from manipulating if he has a large value. To formulate our measure, consider a mechanism M and such a distribution f ; let t be the corresponding truthful strategy profile, and let s^* be the corresponding BNE strategy profile. We define the incentives of the mechanism M given f as

$$\begin{aligned} \text{INCENTIVES}_{M,f}(t,s^*) &= \left\| \left\| t_i(v_i) - s^*_i(v_i) \right\|_2 \right\|_2^f \quad (6) \\ &= \sum_{i=1}^n \sqrt{\int_{v_i} f_i(v_i) \cdot (\|t_i(v_i) - s^*_i(v_i)\|_2)^2 dv_i}. \end{aligned}$$

To provide intuition for this measure, we explain the components of Equation (6). The innermost L_2 metric captures the distance between a bidder i 's truthful bid and the BNE bid for a specific valuation. Next, $\|\cdot\|_2^f$ is a standard continuous weighted L_2 -norm to aggregate over all possible auction instances (i.e., valuations). Finally, the outer L_1 -norm simply aggregates over all bidders.

6. Results for LLG

In this section, we study the LLG domain and several novel variants. We first focus on LLG for several reasons: First, there are existing theoretical results for some rules that provide a benchmark for our experiments. Second, solving for the BNE gets exponentially harder as the domain gets more complex. LLG is simple enough that we can solve for the BNE strategies for a large number of rules with high precision. That said, as we will show, both the design space and the resulting BNE structure are surprisingly subtle and intricate. Thus, it is important to understand the BNEs of a small domain before moving to a larger one.

6.1. LLG UNIFORM

In LLG, there are two items, A and B . There are two local bidders, each only interested in item A or B , respectively, and one global bidder who wants both items simultaneously. Prior work has focused on LLG UNIFORM, where each bidder's value is drawn independently, with local bidders' values drawn from $U[0,1]$ and the global bidder's value drawn from $U[0,2]$.

BNEs of MRC-selecting payment rules are complex to study analytically, and existing theoretical results are only available for LLG. Prior work has shown that the BNE strategies of the local bidders require an *additive* shading in this setting.¹³

Proposition 1 (Goeree and Lien 2016, Ausubel and Baranov 2020a). *In LLG UNIFORM, a Bayes–Nash equilibrium of the QUADRATIC rule is for the global bidder to be*

truthful and for the local bidders to bid: $\hat{v} = \max(0, v - (3 - 2\sqrt{2})) \approx \max(0, v - 0.17)$.

In Table 1, we provide the corresponding results from our computational search approach.¹⁴ The way to read this (and all following tables) is as follows. The QUADRATIC results are always provided in the first row (and here, they correspond to the BNE in Proposition 1). For QUADRATIC, we report the absolute values of the measures from Section 5. Recall that EFFICIENCY measures the fraction of the maximum social welfare achieved by the rule in BNE, INCENTIVES is a measure for how far away from truthful the BNE is, and REVENUE measures the fraction of the VCG revenue that the rule achieves in BNE. Thus, in all settings, VCG achieves 100% efficiency, 0 incentives, and 100% revenue. At first sight, it may look like VCG should be preferable over QUADRATIC. However, recall that VCG is not an MRC-selecting rule. It suffers from the problems we discussed in Section 1.1 (in particular, the *possibility* of very low, noncompetitive levels of revenue).

The next three rows of Table 1 show the top rules by each dimension (efficiency, incentives, and revenue). The individual entries for these rules represent the *multiplicative improvement* for each measure relative to QUADRATIC. In this case, FRACTIONAL*($R = p_{\text{SHAPLEY}}, W = \pi_{\text{VCG}}, A = 10$) is best by both efficiency and revenue, providing evidence that the Shapley value can be useful in the design of MRC-selecting rules. The rule is able to increase the efficiency (relative to QUADRATIC) by 0.29% and revenue by 1.43%. In this setting, the best rule by incentives is the FRACTIONAL*($R = p_{\text{SHAPLEY}}, W = \pi_{\text{SHAPLEY}}, A = 5$) rule, where the reference point is also Shapley and where an amplified version of the Shapley payoff is used for weighting. By contrast, in terms of incentives, the non-MRC first-price rule is worse than QUADRATIC by -412.6% (see Online Supplement G).

6.2. LLG with Correlation

We now expand on the basic LLG structure, by introducing correlation in the bidders' values: instead of drawing the values independently, we now draw their values from a joint distribution. We use *copulae* to define these distributions, a method that lets us separate the specification of the *marginal* distributions through which each bidder views its distribution in isolation from the *coupling*, which describes the joint structure among these marginals.

Formally, Sklar's theorem (Sklar 1959) states that *all* multivariate cumulative distribution functions (CDFs) $F(x_1, \dots, x_d) = \mathbb{P}(X_1 \leq x_1, \dots, X_d \leq x_d)$ can be represented as $F(x_1, \dots, x_d) = C(M_1(x_1), \dots, M_d(x_d))$, where the M_i are the marginal CDFs in each of d dimensions (e.g., $M_i(x) = \mathbb{P}(X_i \leq x)$), and C is a *copula* (which is a joint CDF with uniform marginals). The theorem also

Table 1. Results for LLG(MD=UNIFORM)

Result	Rule QUADRATIC	Efficiency (%)	Incentives	Revenue (%)	Avg. (%)
		98.03	0.32	91.30	
		Improvement over QUADRATIC (%)			
Best efficiency	FRACTIONAL*(R = p_{SHAPLEY} , W = π_{VCG} , A = 10)	0.29	2.69	1.43	1.47
Best incentives	FRACTIONAL*(R = p_{SHAPLEY}^M , W = p_{SHAPLEY}^{-1} , A = 3)	0.26	5.10	1.24	2.20
Best revenue	FRACTIONAL*(R = p_{SHAPLEY} , W = π_{VCG} , A = 10)	0.29	2.69	1.43	1.47

Notes. The first row shows the performance of QUADRATIC. The subsequent rows show the top rules for each dimension.

provides that C will be unique if the F_i are continuous. The converse of the theorem lets us create multidimensional models by combining marginal distributions M_i with a copula C to create a joint distribution $C(M_1(x_1), \dots, M_d(x_d))$. First, we consider several choices for C ; in Section 6.3 we consider choices for M_i , and in Section 6.4 we then consider the cross product of these choices.

To model correlation, we adopt standard Gaussian copulae, which use a multivariate normal distribution for the coupling function. We consider two correlation structures: (a) SAME, which establishes correlation between both bidders interested in the *same* item (i.e., between a given local bidder and the global bidder), and (b) CROSS, which establishes correlation between the local bidders.¹⁵ In both cases, the correlation constant is 0.5.

Results for SAME-side correlation are provided in Table 2. The third row of the table illustrates that a rule that is very good by one dimension may be worse on others. We will seek to address this in Section 6.5 by finding good all-rounder rules. See Online Supplements I and J for results on SAME-side and CROSS-side correlation, where in both cases, we also vary the intensity of the correlation.

6.3. LLG with BETA Marginals (Uncorrelated)

We employ a BETA distribution for our marginals, as it approximates the shape of many familiar distributions with just two parameters.¹⁶ We use five parameterizations of the BETA distribution for the local bidders' values (see Figure 2 in Online Supplement F). Once one employs a skewed distribution, the relative bidder strength between the local and global bidders may no longer match that of the UNIFORM case (i.e., the means

of the bidders' value distributions will be in a different ratio to each other). To address this, we linearly *calibrate* the distributions (unless explicitly mentioned) to ensure that the ratio of the means of the bidders' value distributions is the same as for UNIFORM. We also experimented with *uncalibrated* settings, which we include in Online Supplements K, L, and M.

Table 3 provides results for LLG(MD = BETA(3,1/3)), which is an example of the types of results we see in settings with skewed marginals. We again observe that a different rule is optimal for each dimension. However, all three rules perform quite well in all dimensions relative to QUADRATIC. Results for the other distributions are provided in Online Supplement K.

6.4. Maximum Improvements Relative to QUADRATIC

Modeling the joint distribution of value among the bidders using a copulae lets us mix and match between various marginal distributions and types of correlation. We have investigated the full cross product of correlations we discussed in Section 6.2 with the set of marginal distributions discussed in Section 6.3. The full set of results is presented in Online Supplements L and M.

We now briefly point toward those rules that achieve the *largest improvement* over QUADRATIC in any single setting. In terms of efficiency, FRACTIONAL*(R = BID, W = π_{SHAPLEY} , A = 2.0) achieves a 3.99% improvement over QUADRATIC in LLG(MD = BETA(3,1/3), CORR = CROSS, UNCALIBRATED); see Table 4. Considering the fact that many large-scale CAs allocate resources worth billions of dollars, an efficiency improvement of this magnitude is very significant. Note that the same rule, in the same setting, also achieves an incentive improvement of 47.20%. This demonstrates that MRC-selecting rules

Table 2. Results for LLG(MD=UNIFORM, CORR=SAME)

Result	Rule QUADRATIC	Efficiency (%)	Incentives	Revenue (%)	Avg. (%)
		98.41	0.22	95.19	
		Improvement over QUADRATIC (%)			
Best efficiency	FRACTIONAL*(R = ZERO, W = π_{VCG} , A = 2)	0.24	6.77	2.65	3.22
Best incentives	FRACTIONAL*(R = ZERO, W = p_{SHAPLEY} , A = 0.5)	0.17	8.46	2.93	3.85
Best revenue	FRACTIONAL*(R = ZERO, W = $\pi_{\text{SHAPLEY}}^{-1}$, A = 1)	0.04	-0.61	4.00	1.14

Notes. The first row shows the performance of QUADRATIC. The subsequent rows show the top rules for each dimension, with the entry inducing selection shaded in grey.

Table 3. Results for LLG(MD=BETA(3,1/3),CORR=CROSS,UNCALIBRATED)

Result	Rule	Efficiency (%)	Incentives	Revenue (%)	
	QUADRATIC	97.79	0.59	88.24	
		Improvement over QUADRATIC (%)			Avg. (%)
Best efficiency	FRACTIONAL*(R = p_{VCG} , W = π_{VCG}^{-1} , A = 3)	1.19	21.00	14.68	12.29
Best incentives	FRACTIONAL*(R = $p_{SHAPLEY}^M$, W = $p_{SHAPLEY}^{-1}$, A = 3)	1.15	25.74	12.28	13.06
Best revenue	FRACTIONAL*(R = BID, W = $p_{SHAPLEY}$, A = 2)	0.99	6.98	15.80	7.92

Notes. The first row shows the performance of QUADRATIC. The subsequent rows show the top rules for each dimension, with the entry inducing selection shaded in grey.

exist for which, in some settings, their equilibrium strategies are significantly closer to truthful than QUADRATIC. This performance is even exceeded by FRACTIONAL* (R = ZERO, W = EQUAL, A = 1) in LLG(MD = UNIFORM, CORR = CROSSLARGE), where this rule achieves an incentive improvement of 48.25% over QUADRATIC (see Table 8 in Online Supplement J). In terms of revenue, FRACTIONAL* (R = BID^M, W = $p_{SHAPLEY}^{-1}$, A = 2) achieves a 23.43% improvement over QUADRATIC in LLG(MD = BETA(3,1/3), CORR = CROSS) (see Table 29 in Online Supplement M).

6.5. Best All-Rounder Rules

In the previous section, we have evaluated our rules one setting at a time. When auctioneers have good information about their setting structure, this enables the selection of very high-performing rules, even if these rules perform poorly elsewhere in the setting space. However, in practice, auctioneers may not know the exact structure of the setting in which they are operating. Accordingly, we now seek good “all-rounder” rules that are widely applicable.

Different auctioneers might place different emphasis on each of our evaluation dimensions. In the absence of such knowledge, we opt to take a simple average over all three dimensions and then rank our rules by this average. Table 5 shows the top 20. We see that the best rule achieves an 8.22% average improvement (across all three dimensions) over QUADRATIC, across all 29 settings.

Seven of the top 20 rules actually beat QUADRATIC in every dimension in every setting; those rules are highlighted in grey in Table 5. The other rules beat QUADRATIC in almost all of the 29 settings (specifically, in 26,

27, or 28 settings). However, we do not consider it to be an exclusion criterion if a rule loses to QUADRATIC in a few settings. In fact, we find that QUADRATIC performs very well in some domains, where it is almost impossible to beat. It would not make sense to restrict our search for good all-rounder rules to only those that beat QUADRATIC everywhere.

Looking at Table 5, we observe that Shapley-based rules are ubiquitous. Indeed, all seven rules that beat QUADRATIC everywhere use Shapley payments as a reference point. Bosshard and Seuken (2021b) have recently performed a theoretical analysis of such rules. They found that using Shapley payments as a reference point typically leads to a low *local manipulability* of the rule, which may explain its good performance (see Online Supplement A for further discussion).

Unfortunately, the Shapley value is #P-hard for many standard games (Deng and Papadimitriou 1994), and we are not aware of polynomial-time algorithms that compute exact Shapley values for our setting. For LLG, this is not a problem, but in larger domains, this computational complexity becomes prohibitively expensive, and consequently, we must omit the Shapley-based rules from our analysis in LLLLGG. In future work, one could try adapting previously proposed approximation algorithms for computing the Shapley value to our framework. For example, Agarwal et al. (2019, algorithm 2) suggest a way to approximate Shapley values based on sampling.

7. Results for LLLLGG

In this section, we take the rules that worked well in the LLG domain and seek to find out if they also perform well in the larger LLLLGG domain introduced

Table 4. Results for LLG(MD=BETA(3,1/3),CORR=CROSS,UNCALIBRATED)

Result	Rule	Efficiency (%)	Incentives	Revenue (%)	
	QUADRATIC	95.00	0.59	143.59	
		Improvement over QUADRATIC (%)			Avg. (%)
Best efficiency	FRACTIONAL*(R = BID, W = $\pi_{SHAPLEY}$, A = 2)	3.99	47.20	10.02	20.40
Best incentives	FRACTIONAL*(R = BID, W = $\pi_{SHAPLEY}$, A = 2)	3.99	47.20	10.02	20.40
Best revenue	FRACTIONAL*(R = BID, W = $\pi_{SHAPLEY}$, A = 2)	3.99	47.20	10.02	20.40

Notes. The first row shows the performance of QUADRATIC. The subsequent rows show the top rules for each dimension, with the entry inducing selection shaded in grey.

Table 5. Results Showing the Top 20 All-Rounder Rules

No.	Rule	Efficiency (%)	Incentives	Revenue (%)	Avg. (%)
		97.71	0.35	63.72	
Best 20 all-rounder rules		Avg. improvement over QUADRATIC (%)			
1	FRACTIONAL*(R = ZERO, W = π_{VCG} , A = 0.5)	0.92	17.17	6.57	8.22
2	FRACTIONAL*(R = ZERO, W = $\pi_{SHAPLEY}$, A = 0.5)	0.86	17.17	5.90	7.98
3	FRACTIONAL*(R = ZERO, W = π_{VCG} , A = 1)	0.85	16.92	5.73	7.83
4	FRACTIONAL*($p_{SHAPLEY}^M$, W = BID^{-1} , A = 3)	0.82	16.74	5.05	7.54
5	FRACTIONAL*(R = p_{VCG} , W = $p_{SHAPLEY}^{-1}$, A = 1)	0.82	16.30	5.34	7.49
6	FRACTIONAL*($p_{SHAPLEY}^M$, W = $p_{SHAPLEY}^{-1}$, A = 3)	0.81	16.60	5.00	7.47
7	FRACTIONAL*(R = ZERO, W = BID , A = 0.5)	0.79	16.23	5.35	7.46
8	FRACTIONAL*(R = $p_{SHAPLEY}^M$, W = $\pi_{SHAPLEY}^{-1}$, A = 3)	0.81	16.59	4.96	7.46
9	FRACTIONAL*(R = p_{VCG} , W = BID^{-1} , A = 1)	0.80	16.10	5.16	7.35
10	FRACTIONAL*(R = p_{VCG} , W = π_{VCG}^{-1} , A = 2)	0.79	15.72	4.94	7.15
11	FRACTIONAL*(R = ZERO, W = $p_{SHAPLEY}$, A = 0.5)	0.75	15.44	5.01	7.07
12	FRACTIONAL*(R = ZERO, W = EQUAL, A = 1)	0.87	13.42	6.82	7.04
13	FRACTIONAL*(R = p_{VCG} , W = $\pi_{SHAPLEY}^{-1}$, A = 1)	0.71	14.56	4.48	6.58
14	FRACTIONAL*(R = $p_{SHAPLEY}^M$, W = π_{VCG}^{-1} , A = 3)	0.66	13.57	3.89	6.04
15	FRACTIONAL*(R = $p_{SHAPLEY}$, W = π_{VCG} , A = 3)	0.65	13.42	3.84	5.97
16	FRACTIONAL*(R = $p_{SHAPLEY}^M$, W = BID^{-1} , A = 2)	0.58	12.15	3.56	5.43
17	FRACTIONAL*(R = $p_{SHAPLEY}^M$, W = $\pi_{SHAPLEY}^{-1}$, A = 2)	0.57	11.96	3.49	5.34
18	FRACTIONAL*(R = $p_{SHAPLEY}^M$, W = $p_{SHAPLEY}^{-1}$, A = 2)	0.56	11.77	3.46	5.27
19	FRACTIONAL*(R = p_{VCG} , W = $p_{SHAPLEY}^{-1}$, A = 0.5)	0.53	11.09	3.27	4.96
20	FRACTIONAL*(R = $p_{SHAPLEY}$, W = $\pi_{SHAPLEY}$, A = 3)	0.55	10.87	3.05	4.82

Notes. The first row is the average performance of QUADRATIC over all 29 domains. The subsequent rows show the top rules by their average improvement over QUADRATIC. Rules that beat QUADRATIC in every dimension in every domain are shaded in grey.

by Bosshard et al. (2017, 2020) as a generalization of LLG. This domain is significantly more complex, but numerical BNEs can just barely be computed for it using a powerful compute cluster.¹⁷ Specifically, the LLLLGG domain has eight items and six bidders, each of whom is interested in two bundles. There are four local bidders, each interested in two (overlapping) bundles of two items. And there are two global bidders each interested in two distinct sets of four items. There are significant symmetries in the domain that reduce the complexity of the strategy space. The strategies of the local bidders can be represented as two two-dimensional (2D) surfaces, and the strategy of the global bidder can be represented as a pair of symmetric 2D surfaces (which unifies their computation).

In LLLLGG(MD = UNIFORM), each local bidder draws his value for each bundle from $U[0, 1]$, whereas the global bidders draw their values from $U[0, 2]$. The results for the QUADRATIC rule in this domain are provided in Table 6 (see Online Supplement N for further results in LLLLGG(MD = UNIFORM)). Because QUADRATIC already achieves an efficiency of 99.7% in this domain, there is not much room for improvement. We therefore also consider versions of LLLLGG modified in ways analogous to what we have done in LLG, which we again refer to as *settings*. However, because even LLLLGG(MD = UNIFORM) requires thousands of core-hours to solve for one high-quality numerical BNE, we focus on a single modified setting. Concretely, we analyze a setting with BETA(3,1/3) marginals for the

local bidders and $U[0, 2]$ for the global bidders (without correlation). We selected this setting because it is relatively simple and corresponds to the LLG setting shown in Table 3, where we observe that the choice of rule has a substantial effect.

Because computing even a single BNE in LLLLGG(MD = BETA(3,1/3)) is extremely costly, we cannot exhaustively test all of the rules in this setting. Instead, we select all of the (non-Shapley-based) top 20 rules we have previously identified in Table 5, leading to a set of seven rules to be tested in LLLLGG(MD = BETA(3,1/3)). The results are shown in Table 7. We see that, as in LLG, all seven rules also have a positive average improvement over QUADRATIC in this setting. Thus, the performance enhancements we observed for these rules in LLG generalize to this much larger and more complex setting. Furthermore, our top all-rounder rule from Table 5, FRACTIONAL*(R = ZERO, W = π_{VCG} , A = 0.5), also performs extremely well in LLLLGG(MD = BETA(3,1/3)), leading to an average improvement over QUADRATIC of 12.9%. Only one rule, FRACTIONAL*(R = ZERO, W = EQUAL, A = 1), performs even better, achieving an average improvement over QUADRATIC of 14.2%.

Table 6. Results for the Quadratic Rule in LLLLGG(MD = UNIFORM)

Rule	Efficiency (%)	Incentives	Revenue (%)
QUADRATIC	99.7	0.900	108.1

Table 7. Results for LLLGG(MD=BETA(3,1/3))

No.	Rule QUADRATIC	Efficiency (%)	Incentives	Revenue (%)	Avg. (%)
		97.1	2.170	137.3	
		Improvement over QUADRATIC (%)			
1	FRACTIONAL*(R = ZERO, W = Equal, A = 1)	2.1	36.4	4.0	14.2
2	FRACTIONAL*(R = ZERO, W = π_{VCG} , A = 0.5)	2.0	32.4	4.4	12.9
3	FRACTIONAL*(R = ZERO, W = π_{VCG} , A = 1)	1.8	28.4	4.4	11.5
4	FRACTIONAL*(R = ZERO, W = BID, A = 0.5)	1.5	25.6	4.4	10.5
5	FRACTIONAL*(R = p_{VCG} , W = BID ⁻¹ , A = 1)	1.1	17.5	3.3	7.3
6	FRACTIONAL*(R = p_{VCG} , W = BID ⁻¹ , A = 0.5)	0.5	7.6	1.7	3.3
7	FRACTIONAL*(R = ZERO, W = BID, A = 1)	0.4	4.7	1.4	2.2
8	Reserve price-weighted	0.0	0.3	-0.1	0.1
9	FRACTIONAL*(R = p_{VCG} , W = BID, A = 5)	-2.8	-42.1	-7.6	-17.5
10	First-price	-2.7	-78.7	-12.7	-31.4

Notes. The first row shows the performance of QUADRATIC. In the subsequent rows we show seven of our best all-rounder rules from LLG (nos. 1–7). We also include the reserve price-weighted rule (no. 8), one of the worst rules we identified in LLG (no. 9), and first-price (no. 10). The relative improvement numbers have a standard error of less than 0.1%.

For comparison, we also include the worst rule from LLG that converges in every setting. This rule, FRACTIONAL*(R = BID, W = p_{VCG} , A = 5), performs poorly in LLLGG(MD=BETA(3,1/3)) as well, with an average improvement of -17.5% over QUADRATIC, again confirming that our observations from the LLG domain generalize well to LLLGG. As a reference, we also include the first-price rule, even though it is not an MRC-selecting rule. Naturally, first-price has bad incentives. But we observe that first-price also has bad efficiency (similar to our worst-performing rule) and the worst revenue. This provides good support for using MRC-selecting rules over first-price.

Finally, we also include the reserve price-weighted rule (see Section 4.4), motivated by the use of a similar rule in the 2019 Canadian spectrum auction (Government of Canada 2019).¹⁸ We observe that the rule performs similarly to QUADRATIC, which can be explained by the similarity of their BNEs. Thus, in this experiment, modifying QUADRATIC with reserve price weights only has a modest effect on efficiency, incentives, and revenue—at least compared with our top rules.¹⁹

To obtain an intuition for how our rules operate in LLLGG(MD = BETA(3,1/3)), we plot the BNE strategies

in Figure 2. The BNE strategy for QUADRATIC is shown in yellow and that of one of our top rules, FRACTIONAL*(R = ZERO, W = EQUAL, A = 1), is shown in green. We observe that our rule induces more truthful bidding than QUADRATIC for both local and global bidders with large values (i.e., the BNE bids are closer to the true value). Additionally, our rule also slightly reduces the global bidders’ incentives to overbid when they have a small value. Figure 3 in Online Supplement O shows the BNE for the worst-performing rule. The main observation we see there is that the worst-performing rule provides worse incentives (i.e., it induces more shading in the reported value) for the global bidders when their value is large compared with QUADRATIC.

8. Analysis and Discussion

So far, we have identified a set of rules that perform well in LLG and have seen that their good performance generalizes well to LLLGG. We now take a step back to evaluate whether we can identify any patterns—that is, whether we can answer the question, *what makes a rule a good rule?*

Figure 2. (Color online) BNE Strategies for QUADRATIC and FRACTIONAL*(R=ZERO,W=EQUAL,A=1) in LLLGG(MD=BETA(3,1/3))


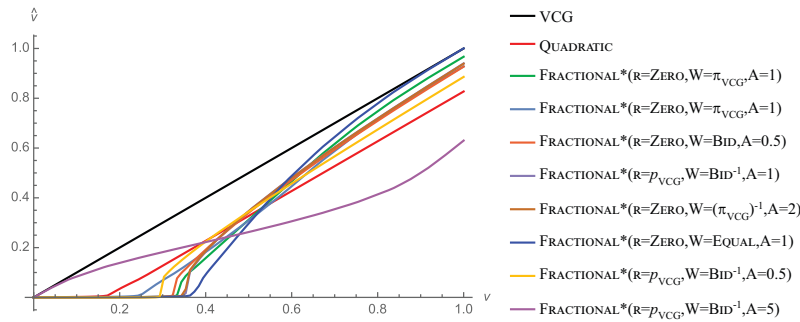


Figure 3. (Color online) BNE Strategies for VCG, QUADRATIC, Seven of Our Best-Performing Rules, and One Poorly Performing Rule in LLG(MD=UNIFORM)

We first take a closer look at the seven non-Shapley rules from Table 5 that we have analyzed in both LLG and LLLGG. To gain some intuition for what incentives these rules provide, we plot the BNE strategies of local bidders under all seven rules (in addition to the worst-performing rule) for LLG(MD = UNIFORM) in Figure 3. We observe a very clear pattern: all of our well-performing rules are large-style rules—that is, they provide much better incentives to bidders with large values (i.e., above 0.5) than does QUADRATIC (by calling them “large-style rules,” we follow the terminology introduced by Parkes (2001)). By contrast, our worst-performing rule is a small-style rule. The differences in incentives also directly translate into differences in utilities.²⁰ We have verified that this pattern (i.e., that our best-performing rules are large-style rules) applies in the other LLG settings as well. Finally, in Figure 2, we have seen that for LLLGG(MD = BETA(3,1/3)), one of our best-performing rules, FRACTIONAL*(R = ZERO, W = EQUAL, A = 1), also provides better incentives to high-valued bidders than does QUADRATIC and thus also behaves similar to a large-style rule in that setting.

A priori, it was not clear that large-style rules would emerge as the best-performing rules.²¹ While a theoretical analysis of this effect is beyond the scope of the present paper, we can provide some intuition for why large-style rules perform well in our analysis. To this end, consider FRACTIONAL*(R = p_{VCG} , W = BID^{-1} , A = 5) in LLG(MD = UNIFORM), which is denoted as rule B in Figure 1. Under this rule, the larger a local bidder’s bid, the closer his payment will be to his VCG payment. If a bidder’s bid is sufficiently large compared with the other bidder’s bid, then the bidder pays essentially his VCG payment (see Figure 1). In that case, a bidder with a large value has negligible incentive to be nontruthful. While no bidder will always face VCG incentives, the larger the bidder’s value, the more likely he is to face (close to) VCG incentives. Thus, a large-valued bidder then optimizes for a distribution of scenarios where he

often faces VCG incentives, such that his optimal strategy is much closer to bidding truthful than for a small-valued bidder.

In LLG(MD = UNIFORM), for a core-selecting rule to have an effect, the local bidders must jointly outbid the global bidder. If none of the bidders has an incentive to overbid (which is the case for all rules in Figure 3), then an efficiency loss must be caused by a local bidder shading his value (such that the two local bidders do not win even though this would be the efficient outcome). Holding the other bid fixed, the smaller a bidder’s shade, the smaller the probability of an efficiency loss occurring. Finally, *conditional on winning*, the value distribution of each local bidder is shifted right (compared with the ex ante uniform distribution). Therefore, the incentives of large-valued bidders matter more for efficiency than the incentives for small-valued bidders. This effect is also nicely exemplified by the worst-performing rule depicted in Figure 3, where small-valued bidders have an incentive to be almost truthful, but large-valued bidders must shade a lot.

In our framework, there are multiple ways to design a large-style rule. One way is exemplified by the top rule in Table 5, FRACTIONAL*(R = ZERO, W = π_{VCG} , A = 0.5). This is an interesting rule, as it uses a zero reference point, which heavily tilts the payments in favor of the large bidders (see Day and Cramton (2012)). Additionally, it uses π_{VCG} as weights (which tilts the payments back toward the small bidders). Finally, it uses a small (0.5) amplification, which deemphasizes the weights, thus making the rule more like QUADRATIC, but not quite. Considering all parameters together, the rule is a dampened version of QUADRATIC with a zero reference point. A rule that is very similar is FRACTIONAL*(R = ZERO, W = BID, A = 0.5) (in row 7), except that it uses BID as the weight instead of π_{VCG} . Given that $\pi_{VCG} = BID - p_{VCG}$, it is intuitive that these two rules perform similarly.

An alternative way to construct a large-style rule is to use a reference point with a more moderate tilting

effect, and instead create the effect via the weights. This is exemplified by the rule $\text{FRACTIONAL}^*(R = p_{\text{VCG}}, W = \text{BID}^{-1}, A = 1)$ in row 9 of Table 5. This rule uses the standard p_{VCG} reference point but combines it with BID^{-1} as weights. The inverted bid weighting has the effect of tilting the payments in favor of the large bidders compared with the unweighted QUADRATIC .²²

We want to emphasize that there is no such thing as an “optimal reference point” or an “optimal weight.” If we consider Table 5 again, we see that among the 20 rules, 4 of our 6 reference points show up and 9 of our 11 weights show up. No clear winner emerges. In fact, our results show that a search for an optimal reference point or an optimal weight is misguided, because it is the *combination* of the reference point, the weights, and the amplification that determines whether a payment rule performs well or not. This is well illustrated by the pair of rules in Table 5 in rows 7 and 9. The rule in row 7, $\text{FRACTIONAL}^*(R = \text{ZERO}, W = \text{BID}, A = 0.5)$, uses BID as weights, whereas the rule in row 9, $\text{FRACTIONAL}^*(R = p_{\text{VCG}}, W = \text{BID}^{-1}, A = 1)$, uses BID^{-1} as weights. Thus, whether the BID should be inverted depends on the reference point. In fact, if we combine the VCG payment reference point with the noninverted BID weight, we obtain a very badly performing rule, $\text{FRACTIONAL}^*(R = p_{\text{VCG}}, W = \text{BID}^{-1}, A = 1)$, whose average improvement over QUADRATIC is -8.88% . If we additionally add an amplification of 5, we obtain our worst-performing rule, $\text{FRACTIONAL}^*(R = p_{\text{VCG}}, W = \text{BID}, A = 5)$, whose average improvement over QUADRATIC is -29.24% . This highlights the importance of choosing the right *combination* of reference point, weights, and amplification.

From our results in LLG and LLLLGG , we have identified seven very good rules (i.e., rules 1–7 in Table 7) that systemically outperform QUADRATIC . We want to highlight two of those rules. First, $\text{FRACTIONAL}^*(R = \text{ZERO}, W = \pi_{\text{VCG}}, A = 0.5)$ stands out, as it was the top rule according to our all-rounder analysis in Section 6 and because it also performed very well in $\text{LLLLGG}(\text{MD} = \text{BETA}(3,1/3))$. Second, $\text{FRACTIONAL}^*(R = \text{ZERO}, W = \text{EQUAL}, A = 1)$ stands out because it is particularly simple and was the top rule in $\text{LLLLGG}(\text{MD} = \text{BETA}(3,1/3))$. Furthermore, this rule had previously been studied by Day and Cramton (2012), albeit only at truth and not in BNE .

When interpreting our results, it is important to note that we have measured the performance of all rules relative to QUADRATIC . Thus, if QUADRATIC performs particularly badly in one dimension, then rules that perform well in this dimension have an advantage. One could also consider choosing another benchmark, which might change our results. However, given that QUADRATIC is currently the most widely used rule in practice, we consider this the most natural benchmark.

One limitation of our approach lies in the *specific* objective we have adopted in our search for the best all-rounder rules (i.e., the average improvement of a

rule compared with QUADRATIC across all dimensions and settings). Of course, other objectives are conceivable, including (1) maximizing the minimum average improvement across all settings, (2) maximizing the minimum improvement across all three dimensions, or (3) maximizing the three dimensions in lexicographic order. One might also add new dimensions such as *fairness*. Our large-style rules shift the benefit to large-valued bidders which may be considered unfair. However, balancing between efficiency and fairness raises challenging new questions (e.g., how to measure fairness). Lubin et al. (2015) already made some progress in this direction, but more work is still needed. Our computational search approach is agnostic to the particular objective of the search, and we do not argue in favor of any one objective. Instead, we have adopted the standard dimensions that have been used in prior work; maximizing the average improvement across the three dimensions was the most straightforward aggregation method. It would be interesting for future work to explore other objectives.

9. Conclusion

We have presented a computational search approach for finding good MRC -selecting payment rules. We have constructed a parameterized design framework to describe any MRC -selecting payment rule, guaranteeing that all rules can be applied in any CA domain. Our results have shown that the *combination* of the reference point, weights, and amplification determines the performance of a rule. Our search identified multiple very good all-rounder rules that beat QUADRATIC on each dimension (efficiency, incentives, and revenue) in almost all settings we have studied. We found that all of these rules are large-style rules (i.e., they provide particularly good incentives to bidders with large values). We have highlighted two particularly promising rules that are simple and outperform QUADRATIC by a significant margin.

Our work illustrates that a computational search approach can be more powerful than designing a mechanism by hand. Designing MRC -selecting rules lends itself to this approach, because the design space can be nicely parameterized and then searched through. It is an interesting topic for future work to explore a computational search approach in other suitable mechanism design domains. Recently, Newman et al. (2020) used a computational search to find optimal parameters for a descending clock auction. Given that, under certain assumptions, the descending clock auction is strategyproof, they assumed straightforward truthful bidding when simulating bidder behavior. However, in many domains, truthful bidding is not a plausible model of bidder behavior, and the computational complexity of equilibrium computation

is likely to be a bottleneck. Therefore, future work on equilibrium computation (be it for auctions or other mechanisms) is important to enable computational search approaches in further domains.

It would be an interesting topic for future work to perform a theoretical analysis of our best-performing rules. Bosshard and Seuken (2021b) have already made some progress regarding an analytical explanation for the advantages of our Shapley-based rules, but there are still many open questions. Given the success of using the Shapley value for the construction of reference points and weights, it would also be interesting to explore other economically well-motivated payoff vectors for this purpose (e.g., the *nucleolus* and the *Nash bargaining solution*).

Going forward, we encourage other researchers to also consider using a computational search approach for mechanism design. Finally, we hope that some of the best rules we have identified may spur new research and that they may be considered for implementation in practice.

Acknowledgments

Authors are listed in alphabetical order. All authors contributed equally to this paper. Some of the ideas presented in this paper were also described in a one-page abstract that was published in the proceedings of the 19th ACM Conference on Economics and Computation (ACM EC'18) (Bünz et al. 2018).

Endnotes

¹ See the comments of the committee for the 2020 Economics Nobel Prize awarded to Milgrom and Wilson at <https://www.nobelprize.org/uploads/2020/09/advanced-economicsciencesprize2020.pdf> (accessed January 16, 2021).

² To distinguish the auctioneer from the bidders, we use “she” or “her” for the auctioneer and “he” or “him” for the bidders.

³ Whether this is indeed a concern depends on the auctioneer’s preferences. If the auctioneer cares about good incentives and *expected* revenue, then VCG may be the best choice (as our results show). However, for applications where even the possibility of a very low-revenue outcome is unacceptable, the auctioneer may deem VCG unsuitable.

⁴ This is in contrast to *automated mechanism design* (Sandholm 2003), where a mechanism is automatically created (by an algorithm) for each specific problem instance (at “runtime,” so to say).

⁵ There are two notable exceptions—namely, the papers by Beck and Ott (2013) and by Bosshard and Seuken (2021a), who both also considered bidders placing bids on items they are not interested in. In general, our framework can straightforwardly capture this as well. In future work, it would be interesting to extend our work in this direction.

⁶ The (weighted) Euclidean distance metric is a natural choice, as it is affected by both the total deviation and the maximum deviation. Furthermore, it guarantees a unique solution to the minimization problem (Day and Cramton 2012). While we could also study other distance metrics within our framework, this is beyond the scope of this paper.

⁷ By choosing reference points and weights that are economically meaningful for any CA instance (such as p_{VCG}), we aim to find rules

whose *performance generalizes* across CA instances. Although we cannot guarantee this a priori, our experiments show that this is indeed the case for our best-performing rules. A by-product of choosing economically meaningful reference points and weights is that one can think about each rule in terms of its components and their properties. This facilitates discussions about the rules and their analysis (similar to QUADRATIC).

⁸ Because reference points and weights are functions of bidders’ value reports, they can be manipulated. This is obvious for BID, but it is also true when using p_{VCG} or π_{VCG} : bidders have control over *other* bidders’ VCG payments, which indirectly gives them some control over their own payment under a rule using p_{VCG} or π_{VCG} as reference points or weights. This motivates the use of alternative reference points and weights that cannot be manipulated (e.g., $ZERO = \vec{0}$).

⁹ In our analysis we also include a recently proposed reserve price-weighted version of QUADRATIC (see Section 4.4).

¹⁰ The total number of rules is fewer than $6 \cdot 11 \cdot 6$, because with EQUAL weights, the amplification has no effect.

¹¹ We adopt the *reserve price* nomenclature to match the existing naming of such weights.

¹² Other measures are also plausible. For example, Ausubel and Baranov (2020a) measured efficiency as $\mathbb{E}_{I-f} \left[\frac{SW_M(I)}{SW_{OPT}(I)} \right]$, which is similar but not identical to Equation (4).

¹³ Ausubel and Baranov (2020a) also provided results for two other rules, with and without correlation.

¹⁴ Values are provided through numerical integration of the BNE strategy over the probability distribution of the setting; because there is no sampling, there is no standard error to report.

¹⁵ Ausubel and Baranov (2020a) considered a form of correlation where the local bidders are either exactly the same or drawn independently. This approach is more amenable to theoretical analysis but is less general and less natural.

¹⁶ Ausubel and Baranov (2020a) considered a distribution similar to $BETA(\alpha, 0)$, but they were limited to Pareto-like shapes.

¹⁷ For all MRC-selecting rules, we compute a 0.01-BNE in LLLLGG; at this level, ϵ is smaller than 0.5% of the maximum value of any bidder. For the first-price rule only, we obtain a 0.015-BNE. The average runtime per rule in LLLLGG at this precision was five days, using a cluster of Intel Xeon E5-650 v4 2.20 GHz processors with 40 logical cores each.

¹⁸ We set a simple per-item reserve price of 0.1 for all items. The effect of the reserve-price weighting is that the global bidders (bidding on four items) are assigned twice the weight of the local bidders (bidding on two items). Therefore, the absolute value of the reserve price does not matter.

¹⁹ One caveat is that we use uniform per-item reserve prices, which constrains the effect of the weighting. It would be interesting future work to explore the properties of this rule with more diverse reserve prices.

²⁰ We have computed the utilities obtained under the different rules for the different value quartiles. The bidders with values between 0.75 and 1 achieve 4.5% more utility under $FRACTIONAL^*(R = ZERO, W = EQUAL, A = 1)$ than under QUADRATIC.

²¹ In fact, Lubin and Parkes (2009) found that small-style rules performed well in a *combinatorial exchange* domain.

²² The fact that there are multiple ways to create large-style rules in our framework makes clear that it does not create a *basis* for MRC-selecting rules. Recall that we built our framework from a carefully chosen set of reference points and weights such that our resulting rules can be applied in any CA domain. If we were working with a basis, then we would instead have a purely computational approach without this property.

References

- Agarwal A, Dahleh M, Sarkar T (2019) A marketplace for data: An algorithmic solution. Karlin A, Immorlica N, Johari R, eds. *Proc. 20th ACM Conf. Econom. Comput.* (ACM, New York), 701–726.
- Ausubel LM, Baranov O (2017) A practical guide to the combinatorial clock auction. *Econom. J. (Lond.)* 127(605):F334–F350.
- Ausubel LM, Baranov O (2020a) Core-selecting auctions with incomplete information. *Internat. J. Game Theory* 49(1):251–273.
- Ausubel LM, Baranov O (2020b) VCG, the core, and assignment stages in auctions. Working paper, University of Maryland, College Park, MD.
- Ausubel LM, Cramton P (2011) Auction design for wind rights. Report to Bureau of Ocean Energy Management, Regulation and Enforcement, U.S. Department of the Interior, Washington, DC.
- Ausubel L, Milgrom P (2002) Ascending auctions with package bidding. *BE J. Theoret. Econom.* 1(1):1–42.
- Ausubel LM, Milgrom P (2006) The lovely but lonely Vickrey auction. Cramton P, Shoham Y, Steinberg R, eds. *Combinatorial Auctions* (MIT Press, Cambridge, MA), 17–40.
- Ausubel L, Cramton P, Milgrom P (2006) The clock-proxy auction: A practical combinatorial auction design. Cramton P, Shoham Y, Steinberg R, eds. *Combinatorial Auctions* (MIT Press, Cambridge, MA), 115–138.
- Balcan M-F, Sandholm T, Vitercik E (2019) Estimating approximate incentive compatibility. Karlin A, Immorlica N, Johari R, eds. *Proc. 20th ACM Conf. Econom. Comput.* (ACM, New York), 867.
- Balseiro S, Kim A, Mahdian M, Mirrokni V (2021) Budget-management strategies in repeated auctions. *Oper. Res.* 69(3):859–876.
- Beck M, Ott M (2013) Incentives for overbidding in minimum-revenue core-selecting auctions. *Proc. Annual Meeting Assoc. Soc. Policy* (Deutsche Zentralbibliothek Leibniz-Informationszentrum für Wirtschaft, Duesseldorf, Germany).
- Bosshard V, Seuken S (2021a) The cost of simple bidding in combinatorial auctions. Biro P, Chawla S, Echenique F, eds. *Proc. 22nd ACM Conf. Econom. Comput.* (ACM, New York), 157.
- Bosshard V, Seuken S (2021b) Shapley-based core-selecting payment rules. Preprint, submitted July 2, <https://arxiv.org/abs/2107.01048>.
- Bosshard V, Bünz B, Lubin B, Seuken S (2017) Computing Bayes-Nash equilibria in combinatorial auctions with continuous value and action spaces. Sierra C, ed. *Proc. 26th Internat. Joint Conf. Artificial Intelligence* (AAAI Press, Palo Alto, CA), 119–127.
- Bosshard V, Bünz B, Lubin B, Seuken S (2020) Computing Bayes-Nash equilibria in combinatorial auctions with verification. *J. Artificial Intelligence Res.* 69:531–570.
- Bünz B, Lubin B, Seuken S (2018) Designing core-selecting payment rules: A computational search approach. Tardos E, Elkind E, Vohra R, eds. *Proc. 19th ACM Conf. Econom. Comput.* (ACM, New York), 109–110.
- Clarke E (1971) Multipart pricing of public goods. *Public Choice* 11(1):17–33.
- Cramton P (2013) Spectrum auction design. *Rev. Indust. Organ.* 42(2):161–190.
- Cramton P, Shoham Y, Steinberg R, eds. (2006) *Combinatorial Auctions* (MIT Press, Cambridge, MA).
- Day RW, Cramton P (2012) Quadratic core-selecting payment rules for combinatorial auctions. *Oper. Res.* 60(3):588–603.
- Day RW, Milgrom P (2008) Core-selecting package auctions. *Internat. J. Game Theory* 36(3):393–407.
- Day R, Milgrom P (2013) Optimal incentives in core-selecting auctions. Neeman Z, Roth A, Vulkan N, eds. *Handbook of Market Design* (Oxford University Press, Oxford, UK), 282–298.
- Day RW, Raghavan S (2007) Fair payments for efficient allocations in public sector combinatorial auctions. *Management Sci.* 53(9):1389–1406.
- Deng Y, Lahaie S (2019) Testing dynamic incentive compatibility in display ad auctions. Teredesai A, Kumar V, Li Y, Rosales R, Terzi E, Karypis G, eds. *Proc. 25th ACM SIGKDD Conf. Knowledge Discovery Data Mining* (ACM, New York), 1616–1624.
- Deng X, Papadimitriou CH (1994) On the complexity of cooperative solution concepts. *Math. Oper. Res.* 19(2):257–266.
- Erdil A, Klemperer P (2010) A new payment rule for core-selecting package auctions. *J. Eur. Econom. Assoc.* 8(2–3):537–547.
- Goeree J, Lien Y (2016) On the impossibility of core-selecting auctions. *Theoret. Econom.* 11(1):41–52.
- Government of Canada (2019) Mathematical formulations for winner and price determination for the combinatorial clock auction in the 600 MHz band. Accessed January 17, 2021, <https://www.ic.gc.ca/eic/site/smt-gst.nsf/eng/sf11449.html>.
- Groves T (1973) Incentives in teams. *Econometrica* 41(4):617–631.
- Klemperer P (2010) The product-mix auction: A new auction design for differentiated goods. *J. Eur. Econom. Assoc.* 8(2–3):526–536.
- Lubin B, Parkes D (2009) Quantifying the strategyproofness of mechanisms via metrics on payoff distributions. McAllester D, ed. *Proc. 25th Conf. Uncertainty Artificial Intelligence* (AUAI Press, Arlington, VA), 349–358.
- Lubin B, Parkes DC (2012) Approximate strategyproofness. *Current Sci.* 103(9):1021–1032.
- Lubin B, Bünz B, Seuken S (2015) New core-selecting payment rules with better fairness and incentive properties. Kominers S, Xia L, eds. *Proc. 3rd Conf. Auctions, Market Mechanisms Appl.* (ACM, New York).
- Marszalec D (2018) Fear not the simplicity—An experimental analysis of auctions for complements. *J. Econom. Behav. Organ.* 152(August):81–97.
- Milgrom P (2007) Package auctions and exchanges. *Econometrica* 75(4):935–965.
- Newman N, Leyton-Brown K, Milgrom P, Segal I (2020) Incentive auction design alternatives: A simulation study. Biro P, Hartline J, Ostrovsky M, eds. *Proc. 21st ACM Conf. Econom. Comput.* (ACM, New York), 603–604.
- Niazadeh R, Hartline J, Immorlica N, Khani MR, Lucier B (2021) Fast core pricing for rich advertising auctions. *Oper. Res.* 70(1):223–240.
- Parkes DC (2001) Iterative combinatorial auctions: Achieving economic and computational efficiency. PhD thesis, University of Pennsylvania, Philadelphia.
- Parkes D (2002) On indirect and direct implementations of core outcomes in combinatorial auctions. Technical report, Harvard University, Cambridge, MA.
- Parkes DC, Kalagnanam J, Eso M (2001) Achieving budget-balance with Vickrey-based payment schemes in exchanges. *Proc. 17th Internat. Joint Conf. Artificial Intelligence*, Vol. 2 (Morgan Kaufmann Publishers, San Francisco), 1161–1168.
- Sandholm T (2003) Automated mechanism design: A new application area for search algorithms. Rossi F, ed. *Proc. Internat. Conf. Principles Practice Constraint Programming* (Springer, Berlin), 19–36.
- Sandholm T (2013) Very-large-scale generalized combinatorial multi-attribute auctions: Lessons from conducting \$60 billion of sourcing. Vulkan N, Roth AE, Neeman Z, eds. *The Handbook of Market Design* (Oxford University Press, Oxford, UK), 379–412.
- Sklar M (1959) Fonctions de répartition à n dimensions et leurs marges. *Publ. Inst. Statist. Univ. Paris* 8:229–231.
- Vickrey W (1961) Counterspeculation, auctions, and competitive sealed tenders. *J. Finance* 16(1):8–37.