On Core-selecting and Core-competitive Mechanisms for Binary Single-Parameter Auctions

Evangelos Markakis Artem Tsikiridis

Dept. of Informatics, Athens University of Economics and Business, Athens, Greece
{markakis, artem}@aueb.gr

Abstract. We focus on the notion of the core, as adjusted in the context of auctions by Ausubel and Milgrom [4]. Core-selecting mechanisms have been known to possess good revenue and fairness guarantees and some of their variants have been used in practice especially for spectrum and other public sector auctions. Despite their popularity, it has also been demonstrated that these auctions are generally non-truthful, except when VCG lies in the core. As a result, current research has focused either on identifying core-selecting mechanisms with minimal incentives to deviate from truth-telling, such as the family of Minimum-Revenue Core-Selecting (MRCS) rules, or on proposing truthful mechanisms whose revenue is competitive against core outcomes. Our results contribute to both of these directions. We start with studying the core polytope in more depth and provide new properties and insights, related to the effects of unilateral deviations from a given profile. We then utilize these properties in two ways. First, we propose a truthful mechanism that is $O(\log n)$-competitive against the MRCS benchmark, which is a quite natural revenue benchmark. Our result is the first deterministic core-competitive mechanism for arbitrary binary single-parameter domains, where only a randomized mechanism was known so far. Second, we study the existence of *non-decreasing* payment rules, meaning that the payment of each bidder is a non-decreasing function of her bid. This property has been advocated by the core-related literature as it implies that the marginal incentive for misreporting is minimized. However, it has remained an open question if there exist MRCS mechanisms that are non-decreasing. We answer the question in the affirmative, by describing a subclass of rules with this property. This can be seen as a further refinement on the set of MRCS rules, towards selecting mechanisms with the most desirable attributes.

1 Introduction

The VCG mechanism has been undoubtedly one of the early landmarks within the field of mechanism design. At the same time however, VCG is rarely preferred in more complex real-life auction scenarios, such as allocation of spectrum or other governmental licences. The shortcomings that have led to this situation have been well summarized by [2], and one of the most prominent drawbacks is the unacceptably low revenue that VCG generates on instances that do not lack competition. The VCG payment corresponds to the externality a bidder imposes on her competitors, and as a result, one can have even zero payments in worst case, giving rise to *free-riders* [3].

To counterbalance this issue, Ausubel and Milgrom [2, 4] adapted the notion of the *core* from the theory of cooperative games and introduced the class of *core-selecting* mechanisms. These mechanisms first select an optimal (welfare-maximizing) allocation as in VCG, but then the payments are set in a way that no coalition of bidders together with the auctioneer can switch to a better outcome, of higher revenue for the auctioneer. It was argued in [2] that a mechanism is of suboptimal performance in terms of revenue precisely when the payments it assigns may *not* be in the core, which is quite common for VCG when the goods exhibit complementarities. Moreover, a non-core outcome can be perceived as unfair by coalitions of bidders, that could be collectively willing to pay more but still were not taken into consideration. Over the last years, core-selecting mechanisms gained even higher support especially among practitioners, due to the fact that they have been successfully implemented for a number of high-profile spectrum auctions, as well as other public sector auctions in several countries [11].

Given the good performance of core-selecting auctions in terms of revenue and fairness, the next natural question is whether we can have strategyproof payment rules in the core. Interestingly, for complement-free settings, VCG can lie in the core. When there are complementarities however, core payments do not generally yield truthful mechanisms [16]. With this negative aspect in mind, research on this topic has focused mainly on two directions. The first direction concerns a game-theoretic analysis of core-selecting mechanisms so as to identify which payments from the core polytope have more desirable incentive properties. As an example of this approach, it has been shown in [12] that selecting a minimum revenue core outcome also minimizes in a certain sense, the total gain from unilateral deviations. When the minimum revenue does not prescribe a unique outcome, a further refinement needs to take place, which is guided again by the incentives to deviate. This has led to the family of quadratic payment rules (see Section 5). In parallel to these results, another way to evaluate such mechanisms is by analyzing the performance of their Bayes Nash equilibria, e.g., [3]. At the moment, the outcomes of these works have not yet led to definite conclusions and there is still an active debate on what are the best core-selecting mechanisms, given also the recent experimental evaluation of [8].

The second direction was initiated by [15] and concerns the design of truthful (hence, not core-selecting) mechanisms whose revenue is competitive against a core outcome. The core benchmark was naturally taken to be the minimum revenue core outcome, given the properties highlighted in the previous paragraph. Hence, a mechanism is then called *α-core-competitive* when it achieves a $1/\alpha$ fraction of the minimum revenue core outcome, for $\alpha \geq 1$. The main results of [15] involved the design of core-competitive mechanisms for a particular single-parameter domain motivated by online ad auctions. For more general combinatorial auctions, one can also obtain core-competitive mechanisms using the results of [24], where a stronger benchmark has been considered. This approach is still worth further investigation, as finding the best ratio against the core benchmark it has remained open for various domains of interest.

Our Contribution. We focus on binary single-parameter domains where each bidder is either accepted or rejected in every outcome. We start in Section 3, with providing new insights and properties on the geometry of the core polytope. Our aim is to understand how the polytope is affected by a unilateral deviation of a bidder from a given profile. To do this, we need to perform a kind of sensitivity analysis for the constraints of the core. In the remaining of the paper, we then make use of the main results of Section 3 in two ways. First, in Section 4, we derive a deterministic $O(\log n)$-core-competitive strategyproof mechanism, where n is the number of bidders. So far, only a randomized mechanism with the same ratio was known, implied by [24]. Our result is the first deterministic core-competitive mechanism for arbitrary single-parameter domains. It also provides a separation between core-competitiveness, and the stronger benchmark of [24], for which an impossibilityt result of $\Omega(n)$ has been known even for single-parameter environments. Second, in Section 5, we focus on the question of identifying more preferred mechanisms among the possible continuum of minimum revenue core-selecting (MRCS) payment rules. This family has been recognized as having better incentive properties among core-selecting mechanisms, and to refine it even further, we study the existence of *non-decreasing* payment rules, meaning that the payment of each bidder is a non-decreasing function of her bid [7, 13]. This property has been advocated, among others, for minimizing the marginal incentive to deviate, but it has remained an open question if there exist MRCS rules satisfying it. We provide a positive answer to this question, by describing a subclass of rules possessing the property, which can be seen as a further refinement towards selecting MRCS mechanisms with the most desirable attributes. Overall, we believe our results shed more light on understanding core-selecting and core-competitive mechanisms, and expect that the properties established here can have even broader appeal and applicability.

1.1 Related Work

The core in the context of auctions was introduced in [2, 4], as a suitable formalism to understand settings where the VCG mechanism underperforms in terms of revenue. In [4], Ausubel and Milgrom also proposed core-selection as a standalone auction design goal by introducing an ascending auction format called the ascending-proxy auction, whose equilibrium outcomes are in the core. The topic soon gained popularity both in theory and in practice, and several follow up works emerged afterwards. A series of important works has focused on exploring different core-selecting Pareto-efficient rules that have minimal incentives to deviate or mechanisms that are core-selecting at equilibrium, see e.g., [10–13, 26, 27] and [3]. The incentives to deviate have been quantified under different metrics and, to our understanding, no final consensus on the most acceptable metric has been reached. Recently, an experimental comparison of Quadratic payment rules [11], was conducted by [8] in an attempt to offer more insights on that front.

Regarding strategyproofness and core-selection, the work of [16] showed that when VCG payments lie in the core, then this is the only truthful mechanism in the core, whereas when VCG is not in the core, there exists no other truthful mechanism that is core-selecting. This reveals a severe incompatibility between truth-telling and core-selection, especially for auction domains that exhibit complementarities. Such domains can arise naturally in spectrum auctions or in auctions related to online advertizing. Nevertheless, efforts have been made to characterize the auction environments where the VCG outcome lies in the core [4, 6, 27, 29]. In [4], it is shown that in domains where the set of feasible allocations form a matroid (e.g. multi-unit auctions [30]), VCG payments always lie in the core, and, therefore, it does not suffer from the shortcomings we have discussed.

In [15], Goel et al. suggested the use of the minimum revenue core-selecting (MRCS) outcome, as a competitive benchmark for the design of truthful mechanisms. In their work, they focus on the so called *Text and Image Ad-Auction*, a special case of knapsack auctions, where k ad slots are being auctioned and each bidder is known to require 1 or k ad slots. They proposed a truthful deterministic mechanism that is $O(\sqrt{\log k})$-core-competitive and a truthful randomized one which is $O(\log\log k)$-core-competitive and these factors are shown to be tight. To our knowledge, this is the only work where a core benchmark has been explicitly used for truthful revenue maximization.

Clearly, the problem of designing truthful mechanisms for maximizing revenue is a fundamental research direction that has attracted considerable attention, especially since the initial works of [14,17,19], see also [21]. Later on, envy-free pricing [20] formed another important approach with several follow up papers. However, these lines of inquiry have mostly focused on environments where goods are substitutes (for which VCG payments are in the core), whereas the core-benchmark is meaningful for environments with complementarities. For such environments, two notable benchmarks have been proposed in [1] for knapsack auctions and in [24] for general combinatorial auctions. We refer the reader to [15] for a detailed comparison of all these benchmarks with the minimum-core-revenue benchmark. The two main takeaways of these comparisons are that, the mechanism of [1] performs arbitrarily bad against the MRCS benchmark, whereas the benchmark of [24] is stronger than MRCS. Hence being α-competitive in the sense of [24], implies being α-core-competitive, for $\alpha \geq 1$. In [24], the authors propose a truthful randomized mechanism for general combinatorial auctions that is $O(\log n)$-competitive against their benchmark and show that this result is tight. Moreover, they complement this finding by showing that no deterministic mechanism can be better that $\Omega(n)$-competitive against their benchmark. Their result implies a randomized $O(\log n)$-core-competitive mechanism for the binary single-parameter setting that we study.

Finally, we stress that by definition the core polytope consists of an exponential number of constraints, which makes its use in mathematical programs challenging. Fortunately, a separation oracle was introduced in [12], but still, each call to the separation oracle requires the solution of a welfare optimization problem. Given these considerations, it is often assumed in the core auction literature that a mechanism has oracle access to a welfare optimization algorithm. In these cases, the complexity measure is the number of oracle calls to the welfare optimization problem. Due to [12], one can deduce then a polynomial upper bound for the number of oracle calls required for the computation of a core point. Obviously, when the underlying welfare optimization can be solved in polynomial time, the mathematical program can also be solved in polynomial time. Recently, in [22] a faster algorithm was presented for computing approximate, Pareto-efficient core payments using only a quasi-linear number of oracle calls. Other algorithms that perform well in practice but admit no runtime guarantees are proposed in [9, 12].

2 Definitions and Preliminaries

2.1 Single-Parameter Domains and Mechanisms

Our work focuses on mechanisms for *binary, single-parameter* domains. We consider a set of bidders $N = \{1, 2, \ldots, n\}$, who can express a request for some type of service (e.g., request for obtaining a set of goods, or access to a facility, etc). Each bidder $i \in N$ has a single private parameter $v_i \geq 0$, which denotes the value derived by bidder i if she is granted the service. The environment is binary in the sense that every bidder will be either accepted or rejected. For every subset $S \subseteq N$, we let $\mathcal{F}(S) \subseteq 2^S$ be the set of feasible allocations for the bidders of S, i.e., the collection of subsets of bidders that can be granted service simultaneously. We assume that $\mathcal{F}(N)$ is *downward-closed*, i.e., for every $X \in \mathcal{F}(N)$ and every $Y \subseteq X$ it holds that $Y \in \mathcal{F}(N)$. We also assume that for every $S \subseteq T$, $\mathcal{F}(S) \subseteq \mathcal{F}(T)$.

An auction mechanism $\mathcal{M} = (X, \mathbf{p})$, in this setting, when run on the set N of agents, consists of an *allocation algorithm* $X : \mathbb{R}^n_+ \mapsto 2^N$ and a *payment rule* $\mathbf{p} : \mathbb{R}^n_+ \mapsto \mathbb{R}^n$. Initially, the auctioneer collects the vector of bids $\mathbf{b} = (b_i)_{i \in N}$, where b_i denotes the bid declared by bidder $i \in N$ (which may differ from v_i). We assume that $b_i \in [0, \infty)$ and that there are no further restrictions on the set of allowed bids. Then, given a bidding profile $\mathbf{b}$, the auctioneer runs the allocation algorithm to determine a feasible allocation $X(\mathbf{b}) \in \mathcal{F}(N)$, and the payment rule to determine the payment vector $\mathbf{p}(\mathbf{b}) = (p_1(\mathbf{b}), \dots, p_n(\mathbf{b}))$, where $p_i(\mathbf{b})$ is the payment requested by bidder i.

We will often need to refer to sub-instances defined by a coalition of bidders. Given a bidding vector $\mathbf{b}$, and a subset of bidders $S \subseteq N$, we denote by $\mathbf{b}_S$ the projection of $\mathbf{b}$ on S, i.e., the vector containing the bids of the members of S. We also denote by $\mathbf{b}_{-i}$ the vector of all bids except for some bidder i. Given a profile $\mathbf{b}$, if we run a mechanism $\mathcal{M} = (X, \mathbf{p})$ on a sub-instance defined by $S \subseteq N$, then $X(\mathbf{b}_S) \in \mathcal{F}(S)$ will denote the resulting allocation and $\mathbf{p}(\mathbf{b}_S)$ will be the corresponding payment vector for the members of S.

We assume that bidders have quasi-linear utilities and hence, given a mechanism $\mathcal{M} = (X, \mathbf{p})$, the final utility of bidder $i \in N$ for a profile $\mathbf{b}$ is $u_i^{\mathcal{M}}(\mathbf{b}) = v_i - p_i(\mathbf{b})$ when $i \in X(\mathbf{b})$ and 0 otherwise (we enforce that losing bidders do not pay anything). We say that $\mathcal{M}$ satisfies individual rationality if for every profile $\mathbf{b}$ and for every bidder $i \in N$ it holds that $u_i^{\mathcal{M}}(\mathbf{b}) \geq 0$. Additionally, a mechanism is truthful, or strategyproof, if for every bidder $i \in N$, every $b_i \geq 0$ and every profile $\mathbf{b}_{-i}$ it holds that $u_i^{\mathcal{M}}(v_i, \mathbf{b}_{-i}) \geq u_i^{\mathcal{M}}(b_i, \mathbf{b}_{-i})$.

Since we are in a single-parameter environment, in order to design truthful mechanisms, we use the characterization of Myerson [25]. In particular, we say that an allocation algorithm X is *monotone* if for every agent $i \in N$ and every profile $\mathbf{b}$, if $i \in X(\mathbf{b})$, then $i \in X(b_i', \mathbf{b}_{-i})$ for $b_i' \geq b_i$. This means that if an agent is selected in an allocation by declaring a bid b_i, then she should also be selected when declaring a higher bid.

Lemma 1. *Given a monotone allocation algorithm X, there is a unique payment rule $\mathbf{p}$ such that $\mathcal{M} = (X, \mathbf{p})$ is an incentive compatible and individually rational mechanism. For every profile $\mathbf{b}$ and every bidder $i \in N$ this payment is given by*

$$p_i(\mathbf{b}) = \inf_{b_i' \in [0, b_i]} \{b_i' : i \in X(b_i', \mathbf{b}_{-i})\}$$

when $i \in X(\mathbf{b})$, and $p_i(\mathbf{b}) = 0$ otherwise.

Lemma 1 is known as Myerson's lemma, and the payments are often referred to as *threshold payments*, since they indicate the threshold below which a bidder stops being selected.

2.2 Welfare Maximization and VCG Payments

For a mechanism $M = (X, \mathbf{p})$, the social welfare produced when run on a profile $\mathbf{b}$ (from the viewpoint of the mechanism since each b_i may differ from v_i) is equal to $\sum_{i \in X(\mathbf{b})} b_i$. Among the most desirable outcomes in mechanism design is to select allocations that achieve maximum welfare. In particular, for a profile $\mathbf{b} \in \mathbb{R}_+^n$, and for any coalition $S \subseteq N$ the *optimal* allocation with respect to $\mathbf{b}_S$ is defined as

$$X^*(\mathbf{b}_S) := \arg\max_{T \in \mathcal{F}(S)} \sum_{i \in T} b_i \tag{1}$$

We will denote by $W(\mathbf{b}_S)$ the maximum social welfare achieved by an optimal allocation. This is also referred to as the *coalitional value* of S: $W(\mathbf{b}_S) := \max_{T \in \mathcal{F}(S)} \sum_{i \in T} b_i = \sum_{i \in X^*(\mathbf{b}_S)} b_i$. When $S = N$, we refer to an optimal allocation by $X^*(\mathbf{b})$ instead of $X^*(\mathbf{b}_N)$, and to the optimal welfare by $W(\mathbf{b})$.

Regarding tie-breaking issues, throughout this work, we assume that a deterministic consistent tie-breaking rule is used to select an allocation, whenever there are multiple optimal allocations at a given profile. For example a fixed ordering on subsets of bidders would suffice to resolve ties.

Fact 1 *Given a bidding vector $\mathbf{b}$, the coalitional value is monotone w.r.t. the set of bidders, i.e. for all $S \subset T \subseteq N$ it holds that $W(\mathbf{b}_S) \leq W(\mathbf{b}_T)$.*

A mechanism is called efficient or welfare-maximizing if for every input profile, it outputs an optimal allocation. The VCG mechanism is the most popular example of an efficient mechanism, where for a bidding profile $\mathbf{b}$, the payment of bidder $i \in X^*(\mathbf{b})$ is the externality she imposes to the other bidders (i.e., the loss to their welfare), defined as

$$p_i^{VCG}(\mathbf{b}) = W(\mathbf{b}_{-i}) - \sum_{j \in X^*(\mathbf{b}) \setminus \{i\}} b_j \tag{2}$$

For every other bidder $i \notin X^*(\mathbf{b})$, we have $p_i^{VCG}(\mathbf{b}) = 0$. For the settings we study, one can easily check that the VCG mechanism is individually rational and strategyproof.

2.3 Core-selecting Payment Rules

The notion of the core as a solution concept originates from cooperative game theory where it captures the fact that coalitions of agents do not have incentives to appeal to a payoff division. To adjust these ideas to the context of auctions, we first define the following quantity, for every coalition $S \subseteq N$ and bidding profile $\mathbf{b}$.

$$\beta(S, \mathbf{b}) := W(\mathbf{b}_S) - \sum_{j \in X^*(\mathbf{b}) \cap S} b_j.$$

This quantity is a generalization of the VCG payment formula, and can be interpreted as the *collective externality* that bidders in $N \setminus S$ impose to the bidders in S. Indeed, with this notation we can restate VCG payments in Equation (2) as $p_i^{VCG}(\mathbf{b}) = \beta(N \setminus \{i\}, \mathbf{b})$, for every bidder $i \in X^*(\mathbf{b})$.

Core-selecting payment rules were initially defined in the space of utility vectors by [4]. In our work we follow the equivalent formulation of [11] that recasts them to the space of payment vectors. For a profile $\mathbf{b}$, the core polyhedron is defined w.r.t. an optimal allocation $X^*(\mathbf{b})$ as follows

$$CORE(\mathbf{b}) = \{\mathbf{p} \in \mathbb{R}^n : \sum_{j \in X^*(\mathbf{b}) \setminus S} p_j \geq \beta(S, \mathbf{b}) \ \forall S \subseteq N, \ p_j = 0 \ \forall j \notin X^*(\mathbf{b})\}. \tag{3}$$

Definition 1. *A payment rule is called core-selecting, if it is individually rational w.r.t. the reported bids, and* $\mathbf{p}(\mathbf{b}) \in CORE(\mathbf{b})$ *for every profile* $\mathbf{b}$*. Furthermore, a mechanism* $\mathcal{M} = (X, \mathbf{p})$ *is a core-selecting mechanism if (i)* $X(\mathbf{b})$ *is a welfare-maximizing allocation for every profile* $\mathbf{b}$*, and (ii)* $\mathbf{p}$ *is a core-selecting payment rule.*

The constraints of the core polytope in (3) require that every coalition of bidders pays at least their collective externality or, in other words, the damage their presence inflicts on the remaining bidders. To provide more intuition, another way to view this is that under a core payment vector, and if bidders are truthful, then every coalition S, together with the auctioneer creates a collective utility at least as high as $W(\mathbf{b}_S)$, which is the best they could achieve if they ran an auction among themselves. In more detail, if u_0 is the auctioneer's utility, which equals $\sum_{j \in N} p_j$, the core constraint for S in (3) is equivalent to:

$$u_0 + \sum_{j \in S} u_j(\mathbf{b}) \geq W(\mathbf{b}_S)$$

Using this formulation, and individual rationality, if the outcome of a mechanism is not in the core, this implies that $u_0 < W(\mathbf{b}_S)$. Hence, there was a coalition that could offer the auctioneer a higher revenue and yet this did not happen.

It is easily verifiable that the pay-your-bid auction, where every winning bidder pays her bid, coupled with the optimal allocation, is a core-selecting mechanism. This rule is sometimes mentioned in the literature as the *seller-optimal* core-selecting payment rule since it maximizes the revenue of the auctioneer with respect to the declared bids. Given that core-selecting mechanisms are not truthful in general, see also [16], a natural quest has been to identify payments in the core where the incentives to misreport are minimized. Formalizing this idea, Day and Milgrom [10] proposed the use of *Pareto-efficient* core payments, which, in the core-literature are also referred to as *bidder-optimal* payment rules.

Definition 2 (Pareto-efficient core payments [10]). *Let* $\mathbf{b}$ *be a bidding profile and* $\mathbf{p} \in CORE(\mathbf{b})$*. We say that* $\mathbf{p}$ *is a Pareto-efficient core payment if for every payment* $\mathbf{p}'$ *such that* $p'_i \leq p_i$ *for every bidder* $i \in X^*(\mathbf{b})$ *with strict inequality for at least one bidder, we have that* $\mathbf{p}' \notin CORE(\mathbf{b})$*.*

A prominent class of Pareto-efficient payment rules in the literature are the Minimum Revenue Core-Selecting (MRCS) rules, i.e., the minimum revenue points in the core, first introduced in [12]. An *MRCS* rule assigns payments given a profile $\mathbf{b}$, that are optimal solutions of the linear program:

$$\min_{\mathbf{p} \in \mathbb{R}^n} \left\{ \sum_{j \in N} p_j : \quad \mathbf{p} \in CORE(\mathbf{b}), \quad \mathbf{p} \leq \mathbf{b} \right\}. \tag{4}$$

It is trivial to check that this is indeed a Pareto-efficient core payment rule. We denote by $\mathsf{MREV}(\mathbf{b})$ the optimal value of the objective function in (4). As shown in [10], the minimum core revenue still gives a better revenue guarantee than VCG, i.e., for a profile $\mathbf{b}$, $\mathsf{MREV}(\mathbf{b}) \geq \sum_{i \in N} p_i^{VCG}(\mathbf{b})$. A further advantage of MRCS rules, established in [12], is that they minimize the *total gains from unilateral deviations*. Finally, it is also interesting to note that whenever the VCG payment belongs to the core, it is the unique MRCS rule [1], bacause it is the unique Pareto-efficient point [10]. Otherwise, the linear program in (4) has a continuum of solutions and a secondary refinement is required in practice to select a particular MRCS payment rule in a disciplined way. We continue this discussion in Section 5, by studying Quadratic Payment Rules, a class of core payment rules which are often used as such a refinement.

2.4 Core-competitive Mechanisms

A different approach has been initiated in [15] concerning revenue guarantees in relation to the core outcomes. Since core-selecting mechanisms are not always truthful (despite their good incentive properties), [15] proposed the design of truthful mechanisms whose revenue is competitive against a core outcome. Given the discussion in Section 2.3, it is quite natural to use as a core benchmark the revenue attained by the MRCS rules. One can evaluate then truthful mechanisms as follows:

Definition 3 ([15]). *Let $\mathcal{M} = (X, \mathbf{p})$ be a truthful mechanism. We say that M is α-core-competitive, with $\alpha \geq 1$, if for any bidding profile $\mathbf{b}$ it assigns a payment vector $\mathbf{p}(\mathbf{b})$ such that*

$$\sum_{i=1}^{n} p_i(\mathbf{b}) \geq \frac{1}{\alpha}\mathsf{MREV}(\mathbf{b})$$

We will follow this approach in Section 4 for single-parameter domains.

3 Insights on the Geometry of the Core

The goal of this section is to establish new insights and properties for the core polytope, and in particular with regard to how the polytope changes when a single bidder declares a higher bid, i.e., we study the relation between $CORE(\mathbf{b})$ and $CORE(b'_i, \mathbf{b}_{-i})$, with $b'_i > b_i$ for some $i \in X^*(\mathbf{b})$. The results we present here will be the key ingredients to prove the two main results of our work in Section 4 and Section 5.

Throughout this section, we assume that for all payment vectors that we consider, we have set $p_j = 0$ for every $j \notin X^*(\mathbf{b})$, for a profile $\mathbf{b}$. We refer the reader to Appendix A for all the missing proofs of this section.

3.1 Warm up: Pareto-efficiency and Individual Rationality within the Core

According to Definition 1, a core-selecting mechanism must be individually rational with respect to the reported bids. In this section, we show that for Pareto-efficient core-selecting payment rules, we have individual rationality for free, and there is, in fact, no need for the auctioneer to explicitly enforce the IR constraints. We start with Lemma 2, which is a straightforward characterization of Pareto-efficient payment rules. It simply says that for every winning bidder i, at least one core constraint that contains the payment of i must be satisfied with equality.

Lemma 2. *Let $\mathbf{b}$ be a bidding profile, and $\mathbf{p} \in CORE(\mathbf{b})$. The vector $\mathbf{p}$ is a Pareto-efficient core payment if and only if for every bidder $i \in X^*(\mathbf{b})$ there exists a coalition $S \subset N$ with $i \notin S$ such that*

$$\sum_{j \in X^*(\mathbf{b}) \setminus S} p_j = \beta(S, \mathbf{b}). \tag{5}$$

We now show that Pareto-efficiency within $CORE(\mathbf{b})$ implies individual rationality with respect to $\mathbf{b}$.

Lemma 3. *A payment rule that for any given profile $\mathbf{b}$ prescribes a Pareto-efficient vector of payments $\mathbf{p} \in CORE(\mathbf{b})$, satisfies $p_i \leq b_i$ for every bidder $i \in X^*(\mathbf{b})$.*

[1] In this case the total gains from unilateral deviations are actually 0, as VCG is an incentive compatible mechanism.

Lemma 3 allows us to omit individual rationality constraints and focus only on the core constraints, when reasoning about Pareto-efficient payment rules. Moreover, using the fact that MRCS payments are Pareto-efficient, we can now simplify the linear program of Equation (4).

Corollary 1. *A payment rule is* MRCS *if, given a profile* $\mathbf{b}$*, it assigns payments that are optimal solutions of the linear program*

$$\min_{\mathbf{p}\in\mathbb{R}^n} \left\{ \sum_{j\in N} p_j : \quad \mathbf{p} \in CORE(\mathbf{b}) \right\}. \tag{6}$$

3.2 The Effects of Unilateral Deviations on the Core

We now aim to understand how the core polytope that forms after a unilateral deviation of a winning bidder is related to the initial core polytope. Initially, we focus on how each of the constraints in the polytope is modified and perform a sensitivity analysis for the term $\beta(S, \mathbf{b})$, the collective externality that appears in the core constraints in (3), for every $S \subseteq N$. Hence, for a given profile $\mathbf{b}$, a bidder $i \in X^*(\mathbf{b})$ and a bid $b_i' > b_i$, we are interested in the relationship between $\beta(S, \mathbf{b})$ and $\beta(S, (b_i', \mathbf{b}_{-i}))$.

To proceed, our analysis will be dependent on the following quantity, defined for an input profile $\mathbf{b}$, a bidder $i \in X^*(\mathbf{b})$, and a coalition $S \subseteq N$ with $i \in S$.

$$t_i(\mathbf{b}_{S\setminus\{i\}}) = \min\{z : \exists T \subseteq S, s.t.\ i \in T \text{ and } \sum_{j\in T\setminus\{i\}} b_j + z = W(z, \mathbf{b}_{S\setminus\{i\}})\} \tag{7}$$

The term $t_i(\mathbf{b}_{S\setminus\{i\}})$ is the minimum bid i should declare to be included in some optimal allocation in an auction where only the bidders from S are present. This is precisely the Myerson threshold payment, for mechanisms where the allocation algorithm produces an optimal allocation when run on input profile $\mathbf{b}_S$. Namely[2], if $i \in X^*(b_S)$, then $t_i(\mathbf{b}_{S\setminus\{i\}}) = p_i^{VCG}(\mathbf{b}_S)$. The following simple lemma can be easily established for the optimal welfare of coalition S.

Lemma 4. *Given a bidding vector* $\mathbf{b}$*, a bidder* $i \in X^*(\mathbf{b})$ *and a bid* b_i' *such that* $0 \le b_i' \le t_i(\mathbf{b}_{S\setminus\{i\}})$*, it holds that*

$$W(b_i', \mathbf{b}_{S\setminus\{i\}}) = W(\mathbf{b}_{S\setminus\{i\}}). \tag{8}$$

The following key lemma encapsulates the effects on the collective externality of S by a unilateral deviation of a bidder $i \in S$.

Lemma 5 (Sensitivity analysis for $\beta(S, \mathbf{b})$). *Let* $\mathbf{b}$ *be a bidding profile. Fix a bidder* $i \in X^*(\mathbf{b})$*, and a coalition* $S \subseteq N$*. Suppose that bidder* i *unilaterally deviates to* $b_i' > b_i$*. Then:*

1. *If* $i \notin S$ *or if* $i \in S$ *and* $b_i \ge t_i(\mathbf{b}_{S\setminus\{i\}})$ *then*

$$\beta(S, (b_i', \mathbf{b}_{-i})) = \beta(S, \mathbf{b}). \tag{9}$$

2. *If* $i \in S$ *and* $b_i < t_i(\mathbf{b}_{S\setminus\{i\}})$ *then*

$$\beta(S, (b_i', \mathbf{b}_{-i})) = \beta(S, \mathbf{b}) - (\min\{b_i', t_i(\mathbf{b}_{S\setminus\{i\}})\} - b_i) \tag{10}$$

Proof. Since the optimal allocation algorithm is monotone and $i \in X^*(\mathbf{b})$, it holds that $X^*(b_i', \mathbf{b}_{-i}) = X^*(\mathbf{b})$, for $b_i' > b_i$. We distinguish the following cases concerning bidder i and the coalition S:

1. $i \notin S$: Then bidder i has no influence on $\beta(S, \mathbf{b})$. By the monotonicity of the allocation algorithm, we have

$$\beta(S, (b_i', \mathbf{b}_{-i})) = W(\mathbf{b}_S) - \sum_{j\in X^*(b_i', \mathbf{b}_{-i})\cap S} b_j = W(\mathbf{b}_S) - \sum_{j\in X^*(\mathbf{b})\cap S} b_j = \beta(S, \mathbf{b}).$$

[2] It can also happen that due to tie-breaking, $X^*(\mathbf{b}_S)$ does not coincide with T from (7), and thus $i \notin X^*(\mathbf{b}_S)$, in which case $t_i(\mathbf{b}_{S\setminus\{i\}}) \neq p_i^{VCG}(\mathbf{b}_S) = 0$.

2. *$i \in S$ and $b_i \geq t_i(\mathbf{b}_{S\setminus\{i\}})$:* By the definition of $t_i(\mathbf{b}_{S\setminus\{i\}})$, we know there exists an optimal allocation $T \in \mathcal{F}(S)$ with respect to $\mathbf{b}_S$, and with $i \in T$. Hence $\sum_{j\in T} b_j = W(\mathbf{b}_S) = \sum_{j\in X^*(\mathbf{b})} b_j$. By the monotonicity of the optimal allocation algorithm, it is true that T is also optimal with respect to $(b'_i, \mathbf{b}_{S\setminus\{i\}})$, for all $b'_i > b_i$. For brevity in the algebraic manipulations below, we denote by X^* the optimal allocation $X^*(\mathbf{b})$. Hence,

$$\begin{aligned}\beta(S,(b'_i,\mathbf{b}_{-i})) &= W(b'_i,\mathbf{b}_{S\setminus\{i\}}) - \sum_{j\in X^*\cap(S\setminus\{i\})} b_j - b'_i = \sum_{j\in T\setminus\{i\}} b_j + b'_i - \sum_{j\in X^*\cap(S\setminus\{i\})} b_j - b'_i \\ &= \sum_{j\in T\setminus\{i\}} b_j + b_i - \sum_{j\in X^*\cap(S\setminus\{i\})} b_j - b_i = W(\mathbf{b}_S) - \sum_{j\in X^*\cap S} b_j = \beta(S,\mathbf{b}).\end{aligned}$$

The second equality holds because $X^*(b'_i,\mathbf{b}_{-i}) = X^*(\mathbf{b})$, and the third equality follows since we argued that T is also optimal for $(b'_i,\mathbf{b}_{S\setminus\{i\}})$.

3. *$i \in S$ and $b_i < t_i(\mathbf{b}_{S\setminus\{i\}})$:* In this case, bidder $i \in X^*(\mathbf{b})$ is not included in any optimal allocation with respect to $\mathbf{b}_S$. We need to consider two subcases. When $b'_i \leq t_i(\mathbf{b}_{S\setminus\{i\}})$ we have:

$$\begin{aligned}\beta(S,(b'_i,\mathbf{b}_{-i})) &= W(b'_i,\mathbf{b}_{S\setminus\{i\}}) - \sum_{j\in X^*\cap(S\setminus\{i\})} b_j - b'_i = W(\mathbf{b}_{S\setminus\{i\}}) - \sum_{j\in X^*\cap(S\setminus\{i\})} b_j - b'_i \\ &= W(\mathbf{b}_S) - \sum_{j\in X^*\cap S} b_j - (b'_i - b_i) = \beta(S,\mathbf{b}) - (b'_i - b_i). \qquad (11)\end{aligned}$$

The second and the third equalities follow from Lemma 4, since both $b'_i \leq t_i(\mathbf{b}_{S\setminus\{i\}})$ and $b_i < t_i(\mathbf{b}_{S\setminus\{i\}})$. In the second subcase, when $b'_i > t_i(\mathbf{b}_{S\setminus\{i\}})$, the unilateral deviation of i enables her to be included in an optimal allocation among bidders in S. Then, we can see that

$$\beta(S,(b'_i,\mathbf{b}_{-i})) = \beta(S,(t_i(\mathbf{b}_{S\setminus\{i\}}),\mathbf{b}_{-i})) = \beta(S,\mathbf{b}) - (t_i(\mathbf{b}_{S\setminus\{i\}}) - b_i).$$

The first equality above follows by applying Equation (9) for the profile $(t_i(\mathbf{b}_{S\setminus\{i\}}),\mathbf{b}_{-i})$, whereas the second equality follows from Equation (11), using $b'_i = t_i(\mathbf{b}_{S\setminus\{i\}})$. Summarizing the two subcases, we obtain $\beta(S,(b'_i,\mathbf{b}_{-i})) = \beta(S,\mathbf{b}) - (\min\{b'_i, t_i(\mathbf{b}_{S\setminus\{i\}})\} - b_i)$, which completes the proof. □

Lemma 5 enables us to prove the two theorems that follow. The first theorem says that for binary single-parameter domains, when a winning bidder declares a higher bid, the space of core payments can only get larger.

Theorem 1. *Let $\mathbf{b}$ be a bidding profile and $i \in X^*(\mathbf{b})$. Then, for every $b'_i > b_i$, $CORE(\mathbf{b}) \subseteq CORE(b'_i,\mathbf{b}_{-i})$.*

Proof. Note first that for $b'_i > b_i$, since the optimal allocation algorithm is monotone and $i \in X^*(\mathbf{b})$, it holds that $X^*(b'_i,\mathbf{b}_{-i}) = X^*(\mathbf{b})$. Consider now a vector $\mathbf{p}$ in $CORE(\mathbf{b})$. We will show that $\mathbf{p}$ is also a member of $CORE(b'_i,\mathbf{b}_{-i})$. This is equivalent to showing that for every $S \subseteq N$, $\mathbf{p}$ satisfies

$$\sum_{j\in X^*(\mathbf{b})\setminus S} p_j \geq \beta(S,(b'_i,\mathbf{b}_{-i})).$$

When $S \subseteq N$ is a coalition such that either $i \notin S$ or $i \in S$ and $b_i \geq t_i(\mathbf{b}_{S\setminus\{i\}})$, then by Lemma 5, we immediately have

$$\sum_{j\in X^*(\mathbf{b})\setminus S} p_j \geq \beta(S,\mathbf{b}) = \beta(S,(b'_i,\mathbf{b}_{-i})).$$

On the other hand, when $i \in S$ and $b_i < t_i(\mathbf{b}_{S\setminus\{i\}})$, then again by Lemma 5 (Equation (10)), and since $\mathbf{p} \in CORE(\mathbf{b})$, we obtain

$$\sum_{j\in X^*(\mathbf{b})\setminus S} p_j \geq \beta(S,\mathbf{b}) = \beta(S,(b'_i,\mathbf{b}_{-i})) + \min\{b'_i, t_i(\mathbf{b}_{S\setminus\{i\}})\} - b_i > \beta(S,(b'_i,\mathbf{b}_{-i})),$$

where the last inequality follows from the facts that $b'_i > b_i$ and $t_i(\mathbf{b}_{S\setminus\{i\}}) > b_i$. □

We note that the set inclusion claimed in Theorem 1 can be strict, i.e., there exists a bidding profile $\mathbf{b}$, a bidder $i \in X^*(\mathbf{b})$ and a $b'_i > b_i$ such that $CORE(\mathbf{b}) \subset CORE(b'_i, \mathbf{b}_{-i})$. An example is included in Appendix A (in the proof of Proposition 1).

The next theorem says that in order to obtain a payment that is in the enlarged polyhedron after a bidder's deviation, the deviating bidder should be charged a payment that exceeds her previous bid.

Theorem 2. *Let $\mathbf{b}$ be a bidding profile and fix a bidder $i \in X^*(\mathbf{b})$. For $b'_i > b_i$, let $\mathbf{p} \in CORE(b'_i, \mathbf{b}_{-i})$ be a payment vector with $p_i \leq b_i$. Then, $\mathbf{p} \in CORE(\mathbf{b})$.*

Theorem 2 will be particularly useful in Section 5.

3.3 A Comment on Revenue Monotonicity of MRCS

Theorem 1 has the following corollary for MRCS core payments, defined in (6).

Corollary 2. *Let $\mathbf{b}$ be a bidding profile. Suppose bidder $i \in X^*(\mathbf{b})$, and let $b'_i > b_i$. Then*

$$\mathsf{MREV}(b'_i, \mathbf{b}_{-i}) \leq \mathsf{MREV}(\mathbf{b}). \tag{12}$$

Proof. Let $\mathbf{p}^* \in CORE(\mathbf{b})$ be an MRCS payment (an optimal solution to the linear program in Equation (6)) for the profile $\mathbf{b}$. Moreover, let $\mathbf{p}' \in CORE(b'_i, \mathbf{b}_{-i})$ be an MRCS solution for the profile $(b'_i, \mathbf{b}_{-i})$. By Theorem 1, it is true that every feasible payment vector $\mathbf{p} \in CORE(\mathbf{b})$ is also in $CORE(b'_i, \mathbf{b}_{-i})$. Therefore, since $\mathbf{p}^* \in CORE(\mathbf{b})$, we have that $\mathbf{p}^* \in CORE(b'_i, \mathbf{b}_{-i})$. Hence,

$$\mathsf{MREV}(\mathbf{b}) = \sum_{j \in N} p^*_j \geq \sum_{j \in N} p'_j = \mathsf{MREV}(b'_i, \mathbf{b}_{-i}).$$

The inequality follows since $\mathbf{p}'$ is an optimal solution for MRCS, a linear program with a minimization objective, for the profile $(b'_i, \mathbf{b}_{-i})$. ☐

Corollary 2 states that a higher willingness to pay by a winning bidder will never lead to an increase of the auctioneer's revenue under MRCS, for all binary single-parameter auctions. This result may look counter-intuitive on a first reading, especially for instances where (12) is satisfied with strict inequality. In the literature, this phenomenon is commonly mentioned as a violation of *revenue-monotonicity*. There are several facets in studying revenue monotonicity, as it concerns the effects on the revenue when adding new bidders, or increasing the offers of the current bidders, or more generally when changing some parameter of the auction. The version we consider here is referred to as *bidder revenue monotonicity* [5].

Pareto-efficient rules that assign payments in the core have been known to be susceptible to violating this property. Namely, it has been shown by [5, 23] that in a multi-parameter domain with at least three items, revenue-monotonicity is violated. Here, we strengthen these results by showing that revenue-monotonicity can be violated in single-parameter auctions as well: we construct instances with single-minded bidders, where a unilateral bid increase by a winning bidder strictly decreases the MRCS revenue.

Proposition 1. *In binary single-parameter auction environments, there exists examples where MRCS rules violate revenue-monotonicity, i.e., Equation* (12) *is satisfied with strict inequality.*

Aside from this discussion, and quite surprisingly, Corollary 2 also plays a crucial role in the analysis of a core-competitive mechanism that we present in Setion 4.

4 An $O(\log n)$-core-competitive Strategyproof Mechanism

In this section, we present a first application of the properties we derived in Section 3. We move away from core-selecting mechanisms with the goal of designing truthful mechanisms that achieve a good revenue approximation with respect to core outcomes. Our main result is a deterministic, truthful mechanism that is also $O(logn)$-core-competitive with respect to the MRCS benchmark. Although we are not analyzing core-selecting mechanisms in this section, the properties of the core, identified in Section 3 (namely Corollary 2 of Theorem 1), will still come in handy for the analysis of our mechanism.

Mechanism $\hat{\mathcal{M}} = (\hat{X}, \hat{\mathbf{p}})$.

Input: A value profile $\mathbf{v} \in \mathbb{R}_+^n$.
Output: An allocation $\hat{X}(\mathbf{v}) \in \mathcal{F}(N)$ and a vector of payments $\hat{\mathbf{p}}(\mathbf{v})$.

1. Find an optimal allocation $X^*(\mathbf{v})$. Let $m = |X^*(\mathbf{v})|$.
2. Let $s_1, \ldots, s_m$ be an ordering of the bidders in $X^*(\mathbf{v})$ such that $v_{s_1} \geq \cdots \geq v_{s_m}$. Let $k \leq m$ be the largest index such that
$$k \cdot v_{s_k} \geq \frac{\mathsf{MREV}(\mathbf{v})}{H_n}. \tag{13}$$
3. Set $\hat{X}(\mathbf{v}) = \{s_1, \ldots, s_k\}$ and $\hat{\mathbf{p}}(\mathbf{v})$ according to Myerson's Lemma (Lemma 1).

Fig. 1. An $O(\log n)$-core-competitive and strategyproof mechanism.

The mechanism is described in Figure 1, where we have used the real valuation profile for the bidders, $\mathbf{b} = \mathbf{v}$ (since we will establish that the mechanism is truthful). We also denote the n-th harmonic number by $H_n = \sum_{i=1}^{n} 1/i = \Theta(\log n)$. In the first step, we find a welfare-maximizing allocation. However, instead of allocating to all bidders in the optimal solution, in the second step the mechanism disqualifies some bidders with values that do not meet a certain cutoff. In case of ties in step 2, it suffices to have a consistent deterministic tie-breaking rule, e.g., given by an ordering on the set of bidders. The mechanism tries, in some sense, to be as inclusive as possible, as long as the value of the last member of $\hat{X}(\mathbf{v})$ is not too small for the coalition to collectively miss the cutoff.

The main result of this section is the following:

Theorem 3. *The mechanism $\hat{\mathcal{M}}$ is individually rational, truthful, and $O(\log n)$-core-competitive.*

Sections 4.2 and 4.3 are devoted to the proof of Theorem 3 and all the missing proofs are in Appendix B. Before proceeding to the proof, we discuss some aspects of the mechanism, and comparisons with other results.

4.1 Remarks on Tightness, Complexity and Other Implications

Our mechanism is applicable to all binary single-parameter auction domains with a downward-closed set of feasible allocations $\mathcal{F}(N)$. In particular, for environments where the VCG payments are not in the core, such as environments that exhibit complementarities or where the set $\mathcal{F}(N)$ is not a matroid, our mechanism is the only known deterministic strategyproof mechanism that is competitive with regard to the MRCS benchmark for arbitrary binary single-parameter domains.

Prior to our work, a randomized, exponential, strategyproof mechanism was known that is also $O(\log n)$-core-competitive [24]. Their result is based on establishing competitiveness against a stronger benchmark, which is the maximum welfare that can be achieved when the highest bidder is ignored. We point the reader to [15] for more detailed comparisons with MRCS. What we find most valuable in obtaining our deterministic matching upper bound is that it yielded a better understanding of the core polytope, through the properties identified in Section 3. On the other hand, the randomized mechanism of [24] does not reveal any further structural properties for the core, since it is centered around a different benchmark. Moreover, our result provides a strict separation on the performance of the two benchmarks, since [24] show that deterministic mechanisms cannot perform better than $\Omega(n)$ for their benchmark even for single-parameter domains. Hence, our mechanism illustrates that the benchmark of [24] is much more stringent, whereas the core benchmark is more amenable to multiplicative approximations and might be more suitable for revenue maximization.

Regarding complexity, our mechanism clearly has a worst-case exponential running time, because it requires the computation of an optimal allocation and of $\mathsf{MREV}(\mathbf{b})$. As already discussed in Section 1.1, the bottleneck of having to solve the welfare maximization problem for various subsets of bidders is not uncommon in the core auction literature, and it is often assumed that the mechanism has oracle access to a welfare maximization algorithm. Given the results derived in [12] for computing $\mathsf{MREV}(\mathbf{b})$, we can conclude that our mechanism can be implemented with a polynomial number of oracle calls to welfare maximization. Faster algorithms have also been proposed for $\mathsf{MREV}(\mathbf{b})$, e.g., [22], but these compute ϵ-bidder-optimal core points and hence are not suitable for our mechanism. Finally, [12] implies that for settings where there exist efficient algorithms for welfare optimization, our mechanism is also implementable in polynomial time. As examples, we mention that this is the case (i) for the Text-and-Image setting of [15], (ii) for Maximum Weight Matching auctions, where bidders represent the edges of a graph and optimal welfare corresponds to a maximum matching[3]. Both of these settings are of interest to us since they possess complementarities, hence there are no truthful mechanisms in the core.

As for tightness, recall that for a given value profile $\mathbf{v}$, the mechanism selects as the set of winning bidders, a subset of the optimal allocation $X^*(\mathbf{v})$. Related to this, for the special case studied in [15], the authors show that for mechanisms that output as an allocation a subset of an optimal allocation, $O(\log n)$-core-competitiveness is the best one can hope for. This implies that our result is tight, and among such mechanisms, it achieves the best possible core-competitiveness .

4.2 Feasibility and Monotonicity of $\hat{X}$

To show that the mechanism always outpiuts a feasible allocation, we use the fact that for a given $\mathbf{v}$, $\hat{X}(\mathbf{v}) \subseteq X^*(\mathbf{v})$. Since the optimal allocation $X^*(\mathbf{v}) \in \mathcal{F}(N)$ and since we have assumed that $\mathcal{F}(N)$ is downward-closed, then $\hat{X}(\mathbf{v})$ is feasible.

Moreover, we claim that the allocation algorithm $\hat{X}$ always outputs a non-empty allocation, i.e., the cutoff set in (13) is always achievable by at least one index $k \in \{1, \ldots, |X^*(\mathbf{v})|\}$. To prove this claim, we define first for a vector of values $v_1 \geq v_2 \geq \cdots \geq v_\ell$, the *maximum uniform price revenue* as $\max_{j\in\{1,\ldots,\ell\}} j \cdot v_j$. The following is a well known lower bound on the uniform price revenue, proposed by [18].

Lemma 6 (Due to [18]). *Given $v_1 \geq \cdots \geq v_\ell$, it holds that* $\max_{j\in\{1,\ldots,\ell\}} j \cdot v_j \geq \dfrac{1}{H_\ell}\displaystyle\sum_{i=1}^{\ell} v_i$.

Using this, we can now prove a lower bound in terms of the MRCS revenue.

Lemma 7. *Let $\mathbf{v}$ be a value profile, and $m = |X^*(\mathbf{v})|$. Let $s_1, s_2, \ldots s_m$ be an ordering of the bidders in $X^*(\mathbf{v})$ by their value in a non-increasing order. Then*

$$\max_{j\in\{1,\ldots,m\}} j \cdot v_{s_j} \geq \frac{\mathsf{MREV}(\mathbf{v})}{H_n}$$

Proof. Let $\mathbf{p} \in CORE(\mathbf{v})$ be a core payment of minimum revenue, i.e. $\sum_{j\in X^*(\mathbf{v})} p_j = \mathsf{MREV}(\mathbf{v})$. By invoking Lemma 6 on the values $v_{s_1} \geq v_{s_2} \geq \cdots \geq v_{s_w}$ we have

$$\max_{j\in\{1,\ldots,m\}} j \cdot v_{s_j} \geq \frac{\sum_{j\in X^*(\mathbf{v})} v_j}{H_m} \geq \frac{\sum_{j\in X^*(\mathbf{v})} v_j}{H_n} \geq \frac{\sum_{j\in X^*(\mathbf{v})} p_j}{H_n} = \frac{\mathsf{MREV}(\mathbf{v})}{H_n}.$$

The last inequality follows from the fact that the family of MRCS payment rules are individually rational. ◻

Lemma 7 directly implies that our proposed mechanism always outputs a non-empty solution, i.e., the cutoff value set in (13) will be satisfied by at least one index.
We now show that the allocation algorithm $\hat{X}$ is monotone. Lemma 8 will be the key to establish this argument, which is in turn based on Corollary 2 from Section 3. Lemma 8 states that when a winning bidder increases her bid, the allocation algorithm $\hat{X}$ may only increase the number of bidders it serves.

Lemma 8. *For every value profile $\mathbf{v}$, bidder $i \in \hat{X}(\mathbf{v})$ and every $v'_i > v_i$ it holds that*

$$|\hat{X}(\mathbf{v})| \leq |\hat{X}(v'_i, \mathbf{v}_{-i})|.$$

Proof. Suppose for contradiction that this is not true, i.e. there exists a profile $\mathbf{v}$ with a bidder $i \in \hat{X}(\mathbf{v})$ and a bid $v'_i > v_i$ for which $|\hat{X}(\mathbf{v})| > |\hat{X}(v'_i, \mathbf{v}_{-i})|$. Since $i \in X^*(\mathbf{v})$ and due to the fact that the welfare-maximizing algorithm is monotone, it holds that $i \in X^*(v'_i, \mathbf{v}_{-i})$ as well. Let $\mathbf{s}$ be the ordering of the players in $X^*(\mathbf{v})$,

[3] These auctions can be motivated by facility location and franchising considerations. The auctioneer can be seen as a company aiming to place stores that should not be on the same neighborhood or on the same street.

produced by the mechanism at step 2, on input $\mathbf{v}$, and let $\mathbf{s}'$ be the corresponding ordering of bidders in $X^*(v_i', \mathbf{v}_{-i})$ on input $(v_i', \mathbf{v}_{-i})$. Let $k = |\hat{X}(\mathbf{v})|$ and $k' = |\hat{X}(v_i', \mathbf{v}_{-i})|$. By our assumption, $k' < k$. Bidder i can only be at a lower index in the ranking $\mathbf{s}'$ compared to her position at $\mathbf{s}$, since she has unilaterally deviated to $v_i' > v_i$. This implies that $v_{s_k'} \geq v_{s_k}$. To verify this, either the bidder at position k in $\mathbf{s}'$ has remained the same but with equal or higher value (in case bidder i is at position k) or bidder i has moved up in the ranking and it has displaced some bidder with a higher value, i.e., with an initial index $s_j < s_k$ to position k. However, this yields

$$k \cdot v_{s_k'} \geq k \cdot v_{s_k} \geq \frac{\mathsf{MREV}(\mathbf{v})}{H_n} \geq \frac{\mathsf{MREV}(v_i', \mathbf{v}_{-i})}{H_n}.$$

The second inequality follows from what we have assumed for the execution of the mechanism on input $\mathbf{v}$, whereas the third inequality follows from Equation (12) of Corollary 2. This means that k bidders can still be served on input $(v_i', \mathbf{v}_{-i})$, and hence k' is not the largest index of bidders who can meet the cutoff of (13) under $(v_i', \mathbf{v}_{-i})$. This is a contradiction and it concludes the proof. ◻

We now prove that the allocation algorithm $\hat{X}$ is monotone.

Lemma 9. *The allocation algorithm $\hat{X}$ is monotone, i.e., given a profile $\mathbf{v}$, for every bidder $i \in \hat{X}(\mathbf{v})$ and every $v_i' > v_i$ it is true that $i \in \hat{X}(v_i', \mathbf{v}_{-i})$.*

Remark 1. Note that we only use the quantity $\mathsf{MREV}(\mathbf{v})$ for determining a cutoff point in step 2 of the mechanism. We do not use the individual MRCS payments that arise by the computation of $\mathsf{MREV}(\mathbf{v})$ (which on their own would not yield a truthful mechanism).

4.3 Payments and Revenue Guarantee

By Lemma 9 the allocation rule $\hat{X}$ is monotone and hence, by Myerson's Lemma, each bidder must pay her threshold price, to obtain a mechanism that is incentive compatible and individually rational in a single-parameter setting. Hence, with regard to the proof of Theorem 3, the only statement we are left to prove is that $\hat{\mathcal{M}} = (\hat{X}, \hat{\mathbf{p}})$ is $O(\log n)$-core-competitive. Lemma 10 provides a relationship that is satisfied by the threshold payment of each winning bidder and that will be crucial to obtain this revenue guarantee.

Lemma 10. *Given a value profile $\mathbf{v}$, the threshold payment $\hat{p}_i(\mathbf{v})$ of every bidder $i \in \hat{X}(\mathbf{v})$ for the mechanism $\hat{\mathcal{M}} = (\hat{X}, \hat{\mathbf{p}})$ satisfies $\hat{p}_i(\mathbf{v}) \geq p_i^{VCG}(\mathbf{v})$ and, additionally,*

$$\hat{p}_i(\mathbf{v}) \geq \frac{\mathsf{MREV}(\hat{p}_i(\mathbf{v}), \mathbf{v}_{-i})}{|\hat{X}(\hat{p}_i(\mathbf{v}), \mathbf{v}_{-i})| \cdot H_n}. \tag{14}$$

Proof. Fix a bidder $i \in \hat{X}(\mathbf{v})$. By definition, her threshold payment $\hat{p}_i(\mathbf{v})$ is the minimum bid $v_i' \leq v_i$ she can unilaterally deviate to, so that $i \in \hat{X}(v_i', \mathbf{v}_{-i})$[4]. Recall that, for every profile $\mathbf{v}$, the first step of the allocation algorithm $\hat{X}$ is to find the optimal allocation $X^*(\mathbf{v})$. Hence, since by Myerson's Lemma, the threshold payments for the algorithm X^* are the VCG payments, we can establish that $\hat{p}_i(\mathbf{v}) \geq p_i^{VCG}(\mathbf{v})$.

To prove (14), consider a bid v_i' which also survives step 2, so that $i \in \hat{X}(v_i', \mathbf{v}_{-i})$. In order for i to be included in an optimal allocation, v_i' must be large enough so that Equation (13) is satisfied for the profile $(v_i', \mathbf{v}_{-i})$. Suppose $v_i' = \hat{p}_i(\mathbf{v})$. Let $\mathbf{s}'$ be the ordering of the bidders produced by step 2 of the mechanism, and let $k = |\hat{X}(\hat{p}_i(\mathbf{v}), \mathbf{v}_{-i})|$. By Equation (13) it holds that

$$k \cdot v_{s_k'} \geq \frac{\mathsf{MREV}(\hat{p}_i(\mathbf{v}), \mathbf{v}_{-i})}{H_n} \Leftrightarrow v_{s_k'} \geq \frac{\mathsf{MREV}(\hat{p}_i(\mathbf{v}), \mathbf{v}_{-i})}{k \cdot H_n}. \tag{15}$$

Additionally, since $i \in \hat{X}(\hat{p}_i(\mathbf{v}), \mathbf{v}_{-i})$, her bid cannot be smaller than the bid of the last bidder included in $\hat{X}(\hat{p}_i(\mathbf{v}), \mathbf{v}_{-i})$, as otherwise she would not win. Therefore, $\hat{p}_i(\mathbf{v}) \geq v_{s_k'}$ and the proof follows by combining this fact with Equation (15). ◻

We can now prove the revenue guarantee and conclude the proof of Theorem 3.

Lemma 11. *The mechanism is $O(logn)$-core-competitive.*

[4] Actually, one needs to be a bit more careful with the analysis since the threshold price is the infimum over the set of bids that make bidder i a winner. This leads to a more detailed technical analysis, which we defer to the full version of this work.

5 A Class of Non-decreasing Quadratic Payment Rules

In this section, we illustrate a second application of our results of Section 3, focusing on an important family of quadratic core-selecting payment rules.

5.1 Quadratic Payment Rules

As we have mentioned in Section 2, when VCG payments are not in the core, the solution space of MRCS payments is always a continuum, in which case the linear program of Equation (4) has infinitely many solutions. Even though, as discussed in Section 2, all these solutions in this face of the core polytope, have been shown in [12] to minimize the gain of deviating, the question remained whether one of these points should be preferred over others and whether there is a disciplined way to single out a solution. This motivated [11, 13] to propose a class of core-selecting mechanisms, based on the idea of picking the point on the minimum revenue face of the core that is the closest in Euclidean distance to a given reference point in the vector space. This payment rule can be expressed using two mathematical programs: the linear program of Equation (6) to compute first $\mathsf{MREV}(\mathbf{b})$, and then a quadratic program, as defined below.

Definition 4. *Let $\mathbf{r} \in \mathbb{R}^n_+$. We call a payment rule $\mathbf{r}$-nearest when, for every vector $\mathbf{b}$, it assigns the payment*

$$\mathbf{p}^{\mathbf{r}}(\mathbf{b}) = \underset{\mathbf{p}\in\mathbb{R}^n}{\arg\min} \left\{ \sum_{j \in X^*(\mathbf{b})} (p_j - r_j)^2 : \quad \mathbf{p} \in CORE(\mathbf{b}), \quad \sum_{j \in X^*(\mathbf{b})} p_j = \mathsf{MREV}(\mathbf{b}) \right\}. \tag{16}$$

In words, the quadratic program of (16) assigns for a bidding profile $\mathbf{b}$, the MRCS payment in the core that is the closest to a given vector $\mathbf{r}$. Alternatively, this quadratic program can be also defined without the MRCS constraint. In this case, it has been shown in [28] that the minimum revenue may not be achieved for certain reference points even for minimization objectives that result to Pareto-efficient payments. In this section, we stick to the version that contains the MRCS constraint. Moreover, since the quadratic program in (16) expresses a minimization of Euclidean distance from a convex set to a fixed point, the following well-known fact is true.

Fact 2 *Given vectors $\mathbf{r}$ and $\mathbf{b}$, the payment vector $\mathbf{p}^{\mathbf{r}}(\mathbf{b})$ is unique.*

A number of vectors have been proposed as the reference point $\mathbf{r}$, for this class of payments. Initially, in [11] Day and Cramton used the VCG payments for a reference point, $\mathbf{r} = \mathbf{p}^{VCG}(\mathbf{b})$, as a refinement of MRCS. The motivation of this choice was the findings of [28] who observed that given a profile $\mathbf{b}$ and a payment $\mathbf{p} \in CORE(\mathbf{b})$ the quantity $p_i - p_i^{VCG}(\mathbf{b})$ represents the bidder's "residual incentive to misreport". Hence, minimizing this quantity (or rather, its square) seemed a reasonable choice within MRCS with good incentives. In parallel to this, Erdil and Klemperer [13], developed a different perspective of what $\mathbf{r}$ should be. They leaned more towards constant payment rules with reference points that do not depend on the bidding profile, as their goal was to minimize *marginal* incentives to deviate. One well-studied and intuitive example is the $\mathbf{0}$-nearest mechanism: pick the point in MRCS that is closest to $\mathbf{0}$. Yet another perspective was given in [3] who have proposed the $\mathbf{b}$-nearest payment rule, i.e., the MRCS payments closest to the actual bid. Overall, quadratic rules form a family of core-selecting mechanisms with many deployments in practice in several countries, especially for spectrum and other public sector auctions [11].

5.2 A Class of Non-Decreasing Quadratic Payment Rules

We now consider the following desirable property for payment rules.

Definition 5. *A payment rule is called non-decreasing, if for every profile $\mathbf{b}$, every bidder $i \in N$ and every $b'_i > b_i$ it holds that*

$$p_i(b'_i, \mathbf{b}_{-i}) \geq p_i(\mathbf{b}). \tag{17}$$

This notion has been defined independently in [13] and [7], with a different motivation in mind. In [13], it is argued that payment rules satisfying this property[5] weakly dominate all other payment rules in terms of the so called *marginal incentive to deviate*. Hence, even though such mechanisms may not be truthful, they possess very desirable incentive guarantees. In [7], another advantage of this property is highlighted, which is of computational nature: limiting our attention to non-decreasing payment rules makes the daunting task of computing Bayes Nash equilibria much simpler.

Hence, it becomes important to understand which mechanisms satisfy this property. It can be seen that the VCG mechanism and the pay-your-bid auction do satisfy (17). In the context of MRCS rules, it is shown in [7], that $\mathbf{p}^{VCG}$-nearest is *not* non-decreasing. To our knowledge, it has remained an open question whether there exist MRCS rules that satisfy (17).

We answer this question in the affirmative, by providing a class of quadratic rules that are non-decreasing. To proceed, given a vector $\mathbf{b}$, for all $i \in N$, define $f_i(b_i)$ to be any non-decreasing function of b_i. Let $\mathbf{f}(\mathbf{b}) = (f_1(b_1), \ldots, f_n(b_n))$. The following is the main result of this section.

Theorem 4. *For every* $\mathbf{f} = (f_1(\cdot), \ldots, f_n(\cdot))$*, where each* $f_i(\cdot)$ *is a non-decreasing function of* b_i*, the* $\mathbf{f}(\mathbf{b})$*-nearest payment rule is non-decreasing for binary single-parameter auction domains.*

Notice that the class of $\mathbf{f}(\mathbf{b})$-nearest rules captures both the well known $\mathbf{0}$-nearest and $\mathbf{b}$-nearest mechanisms that were advocated by [13] and [3] respectively.

Corollary 3. *The* $\mathbf{0}$*-nearest and the* $\mathbf{b}$*-nearest payment rules are both non-decreasing for binary single-parameter auction domains.*

The proof of Theorem 4, is based on both Theorem 1 and Theorem 2 from Section 3. We present it in Appendix C.

6 Conclusions and Future Research

We believe our results shed more light on understanding core-selecting and core-competitive mechanisms, and expect that some of the properties established here (e.g. the properties of the core polyhedron in Section 3, or the analysis of quadratic rules in Section 5) can be of independent interest, and with a broader applicability.

There are still several interesting avenues for further investigation. Regarding core-selecting mechanisms, the recent experimental evaluation of [8] has fueled even more the debate of identifying the most appropriate MRCS rules. We find the notion of non-decreasing payment rules, defined in Section 5, to be a suitable refinement of MRCS towards this direction. It would therefore be interesting to better understand or even to characterize which MRCS rules can satisfy this property.

Regarding the design of truthful mechanisms that are core-competitive, the literature is still very scarce on this, and it would be very desirable to identify special cases of single parameter domains, where one can achieve better than $O(\log n)$-competitiveness, such as the Text and Image setting of [15]. Generalizations to multi-parameter domains would also be enlightening.

References

1. Gagan Aggarwal and Jason D Hartline. Knapsack auctions. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1083–1092. Society for Industrial and Applied Mathematics, 2006.
2. Lawrence Ausubel and Paul Milgrom. The lovely but lonely vickrey auction. P. Cramton, Y. Shoham, R. Steinberg, eds., Combinatorial Auctions, 2006.
3. Lawrence M Ausubel and Oleg V Baranov. Core-selecting auctions with incomplete information. *University of Maryland*, 121, 2010.
4. Lawrence M Ausubel and Paul R Milgrom. Ascending auctions with package bidding. *Advances in Theoretical Economics*, 1(1), 2002.
5. Marissa Beck and Marion Ott. Revenue monotonicity in core-selecting package auctions. *Working paper, Technical Report*, 2009.

[5] Erdil and Klemperer originally called these rules monotonic.

6. Sushil Bikhchandani and Joseph M Ostroy. The package assignment model. *Journal of Economic theory*, 107(2):377–406, 2002.
7. Vitor Bosshard, Ye Wang, and Sven Seuken. Non-decreasing payment rules for combinatorial auctions. In *Proceedings of the Twenty Seventh International Joint Conference on Artificial Intelligence*, pages 105–113, 2018.
8. Benedikt Bünz, Benjamin Lubin, and Sven Seuken. Designing core-selecting payment rules: A computational search approach. In *Proceedings of the 19th ACM Conference on Economics and Computation, (EC 2018)*, page 109, 2018.
9. Benedikt Bünz, Sven Seuken, and Benjamin Lubin. A faster core constraint generation algorithm for combinatorial auctions. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
10. Robert Day and Paul Milgrom. Core-selecting package auctions. *International Journal of Game Theory*, 36(3-4):393–407, 2008.
11. Robert W Day and Peter Cramton. Quadratic core-selecting payment rules for combinatorial auctions. *Operations Research*, 60(3):588–603, 2012.
12. Robert W Day and Subramanian Raghavan. Fair payments for efficient allocations in public sector combinatorial auctions. *Management Science*, 53(9):1389–1406, 2007.
13. Aytek Erdil and Paul Klemperer. A new payment rule for core-selecting package auctions. *Journal of the European Economic Association*, 8(2-3):537–547, 2010.
14. Amos Fiat, Andrew V Goldberg, Jason D Hartline, and Anna R Karlin. Competitive generalized auctions. In *Proceedings of the Thiry-fourth Annual ACM Symposium on Theory of Computing*, pages 72–81. ACM, 2002.
15. Gagan Goel, Mohammad Reza Khani, and Renato Paes Leme. Core-competitive auctions. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 149–166. ACM, 2015.
16. Jacob K Goeree and Yuanchuan Lien. On the impossibility of core-selecting auctions. *Theoretical economics*, 11(1):41–52, 2016.
17. Andrew V Goldberg and Jason D Hartline. Competitiveness via consensus. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 215–222. Society for Industrial and Applied Mathematics, 2003.
18. Andrew V Goldberg, Jason D Hartline, Anna R Karlin, Michael Saks, and Andrew Wright. Competitive auctions. *Games and Economic Behavior*, 55(2):242–269, 2006.
19. Andrew V Goldberg, Jason D Hartline, and Andrew Wright. Competitive auctions and digital goods. In *Proceedings of the Twelfth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 735–744. Society for Industrial and Applied Mathematics, 2001.
20. Venkatesan Guruswami, Jason D Hartline, Anna R Karlin, David Kempe, Claire Kenyon, and Frank McSherry. On profit-maximizing envy-free pricing. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1164–1173. Society for Industrial and Applied Mathematics, 2005.
21. Jason Hartline. Mechanism design and approximation. *Book draft. October*, 122, 2013.
22. Jason Hartline, Nicole Immorlica, Mohammad Reza Khani, Brendan Lucier, and Rad Niazadeh. Fast core pricing for rich advertising auctions. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 111–112. ACM, 2018.
23. Laurent Lamy. Core-selecting package auctions: a comment on revenue-monotonicity. *International Journal of Game Theory*, 39(3):503–510, 2010.
24. Silvio Micali and Paul Valiant. Collusion-resilient revenue in combinatorial auctions. Technical report, Technical Report MIT-CSAIL-TR-2007-052, MIT-CSAIL, 2007.
25. Roger B Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
26. Marion Ott and Marissa Beck. Incentives for overbidding in minimum-revenue core-selecting auctions. Technical report, Working paper, Stanford University, 2013.
27. David C Parkes. On indirect and direct implementations of core outcomes in combinatorial auctions. *Harvard University Technical Report*, 2002.
28. David C Parkes, Jayant R Kalagnanam, and Marta Eso. Achieving budget-balance with vickrey-based payment schemes in exchanges. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence*, pages 1161–1168, 2001.
29. Ryuji Sano. Incentives in core-selecting auctions with single-minded bidders. *Games and Economic Behavior*, 72(2):602–606, 2011.
30. William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961.

A Missing Proofs from Section 3

Proof of Lemma 2.
($\Rightarrow$) For every bidder $i \in X^*(\mathbf{b})$, let $S_i \subset N$ be a coalition that satisfies (5) and does not include bidder i. Suppose for contradiction that $\mathbf{p}$ is not Pareto-efficient. Then, there exists a bidder $k \in X^*(\mathbf{b})$ and a $p'_k < p_k$ such that $(p'_k, \mathbf{p}_{-k}) \in CORE(\mathbf{b})$. However, this vector of payments cannot be feasible since

$$\sum_{j \in X^*(\mathbf{b}) \setminus (S_k \cup \{k\})} p_j + p'_k < \sum_{j \in X^*(\mathbf{b}) \setminus S_k} p_j = \beta(S_k, \mathbf{b}),$$

which is a violation of the core constraint in (3) for the coalition S_k. This implies that $(p'_k, \mathbf{p}_{-k}) \notin CORE(\mathbf{b})$, a contradiction.

($\Leftarrow$) Suppose for contradiction that there exists a bidder $k \in X^*(\mathbf{b})$ such that for every coalition $S \subset N$, that does not include k, Equation (5) does not hold, i.e.

$$\sum_{j \in X^*(\mathbf{b}) \setminus S} p_j > \beta(S, \mathbf{b}).$$

Then, there exists a $p'_k < p_k$ such that the payment $(p'_k, \mathbf{p}_{-k})$ satisfies the core constraint for every coalition $S \subset N$ that does not include bidder k. This implies that $(p'_k, \mathbf{p}_{-k}) \in CORE(\mathbf{b})$, (since the remaining core constraints for coalitions that contain k are satisfied by the fact that $\mathbf{p}$ belongs to the core). But this means that $\mathbf{p}$ is not Pareto-efficient, a contradiction. □

Proof of Lemma 3.
Given a vector $\mathbf{b}$, let $\mathbf{p} \in CORE(\mathbf{b})$ be a Pareto-efficient payment. Fix a bidder $i \in X^*(\mathbf{b})$. Since $\mathbf{p}$ is Pareto-efficient, by Lemma 2 there exists a coalition $S \subseteq N$ that does not include i for which $\sum_{j \in X^*(\mathbf{b}) \setminus S} p_j = \beta(S, \mathbf{b})$.

We distinguish the following cases:

1. $S = N \setminus \{i\}$. In this case, bidder i is asked to pay precisely her VCG payment since

$$\sum_{j \in X^*(\mathbf{b}) \setminus (N \setminus \{i\})} p_j = p_i = \beta(N \setminus \{i\}, \mathbf{b}) = p_i^{VCG}(\mathbf{b}) \le b_i$$

and the last inequality holds since VCG is an individually rational mechanism, due to Fact 1.

2. $S \subset N \setminus \{i\}$. Consider the coalition $S \cup \{i\} \subset N$. Since $\mathbf{p} \in CORE(\mathbf{b})$, by (3) we have

$$\begin{aligned}\sum_{j \in X^*(\mathbf{b}) \setminus (S \cup \{i\})} p_j \ge \beta(S \cup \{i\}, \mathbf{b}) &= W(\mathbf{b}_{S \cup \{i\}}) - \sum_{j \in X^*(\mathbf{b}) \cap (S \cup \{i\})} b_j \ge W(\mathbf{b}_S) - \sum_{j \in X^*(\mathbf{b}) \cap S} b_j - b_i \\ &= \beta(S, \mathbf{b}) - b_i = \sum_{j \in X^*(\mathbf{b}) \setminus S} p_j - b_i.\end{aligned}$$

The second inequality follows from Fact 1 and the last equality from the fact that S satisfies Equation (5) by assumption. By rearranging terms we obtain that $p_i \le b_i$. □

Proof of Lemma 4.
When $b'_i < t_i(\mathbf{b}_{S \setminus \{i\}})$, by the definition of $t_i(\mathbf{b}_{S \setminus \{i\}})$ in Equation 7, bidder i is not included in any optimal allocation when only bidders in S are present. Therefore, since i does not generate any value to the coalition S, her existence might as well be ignored and equation (8) holds.

When $b'_i = t_i(\mathbf{b}_{S \setminus \{i\}})$, even though by the definition of $\bar{b}_i(S)$ bidder i *is* included in an optimal allocation for the auction among bidders in S, we claim that at the same time there exists another optimal allocation when only bidders in S are present that does not include i.

Suppose for contradiction that this is not the case. This means that bidder i belongs to *all* optimal allocations among bidders in S when bidding $t_i(\mathbf{b}_{S \setminus \{i\}})$ against the bids $\mathbf{b}_{S \setminus \{i\}}$. Then, bidder i can bid

$t_i(\mathbf{b}_{S\setminus\{i\}}) - \epsilon$ for a sufficiently small $\epsilon > 0$ and remain a member of all optimal allocations among bidders in S. This, however, is a contradiction, since $t_i(\mathbf{b}_{S\setminus\{i\}})$ is defined as the minimum bid i can issue and be a part of one optimal allocation for set S.

Therefore, since there exists an optimal allocation among bidders in S without i for the profile $(t_i(\mathbf{b}_{S\setminus\{i\}}), \mathbf{b}_{S\setminus\{i\}})$, the coalitional value $W(t_i(\mathbf{b}_{S\setminus\{i\}}), \mathbf{b}_{S\setminus\{i\}})$ can be achieved by the bidders of this particular optimal allocation that does not include i. Hence bidder i can be ignored and equation (8) still holds. □

Proof of Theorem 2. Suppose for contradiction that this is not true, i.e. for a deviating bidder $i \in X^*(\mathbf{b})$ there exists a payment profile $\mathbf{p} \in CORE(b_i', \mathbf{b}_{-i})$ with $p_i \le b_i$ such that $\mathbf{p} \notin CORE(\mathbf{b})$. This implies that there exists a coalition $S \subseteq N$ for which the constraint in (3) is violated, i.e.

$$\sum_{j \in X^*(\mathbf{b})\setminus S} p_j < \beta(S, \mathbf{b}). \tag{18}$$

If S is a coalition with $i \notin S$ or $i \in S$ but with $b_i \ge t_i(\mathbf{b}_{S\setminus\{i\}})$, since $\mathbf{p} \in CORE(b_i', \mathbf{b}_{-i})$, by equation (3) we obtain

$$\sum_{j \in X^*(\mathbf{b})\setminus S} p_j \ge \beta(S, (b_i', \mathbf{b}_{-i})) = \beta(S, \mathbf{b}),$$

where the last equality is due to (9) of the Key Lemma. However, this contradicts inequality (18).
Consider now the case when coalition S is such that $i \in S$ and $b_i < t_i(\mathbf{b}_{S\setminus\{i\}})$. Note that this implies that S cannot be the singleton coalition $\{i\}$, as the minimum bid i must bid to be in an optimal allocation on her own is 0, and we would have $b_i < 0$. Consider the constraint $S \setminus \{i\}$. We have

$$\begin{aligned}
\sum_{j \in X^*(\mathbf{b})\setminus S} p_j + p_i &\ge \beta(S \setminus \{i\}, (b_i', \mathbf{b}_{-i})) \\
&= W(\mathbf{b}_{S\setminus\{i\}}) - \sum_{j \in X^*(\mathbf{b}) \cap (S\setminus\{i\})} b_j \\
&= W(\mathbf{b}_S) - \sum_{j \in X^*(\mathbf{b}) \cap S} b_j + b_i = \beta(S, \mathbf{b}) + b_i. \qquad (19)
\end{aligned}$$

The inequality follows from applying (3) for the constraint $S \setminus \{i\}$ and the second equality from Lemma 4, since $b_i < t_i(\mathbf{b}_{S\setminus\{i\}})$. By combining inequalities (18) and (19) we obtain

$$\beta(S, \mathbf{b}) + b_i < p_i + \beta(S, \mathbf{b}),$$

which is a contradiction. □

Proof of Proposition 1. Consider the following combinatorial auction with 6 single-minded bidders and 3 items for sale, $M = \{A, B, C\}$. For $i = 1, \ldots, 6$ we denote by $b_i(T)$ the bid of i for the set of items $T \subseteq \{A, B, C\}$. Since bidders are single-minded, each bidder declares a single bid of this form. The bids and demands are summarized below:

$$\begin{aligned}
&b_1(\{A\}) = 9.5 &\qquad& b_4(\{A, B\}) = 15 \\
&b_2(\{B\}) = 6 && b_5(\{A, C\}) = 15 \\
&b_3(\{C\}) = 6 && b_6(\{B, C\}) = 11
\end{aligned}$$

An allocation is feasible when each of the three items is assigned to a unique bidder. For instance, bidders 1 and 4 cannot be a part of a feasible allocation. For the vector $\mathbf{b}$ above, it is easy to see that the welfare-maximizing allocation algorithm assigns $\{A, B, C\}$ to bidders $X^*(\mathbf{b}) = \{1, 2, 3\}$. Per Equation (6), the MRCS linear program with variables p_1, p_2, p_3 is:

$$\begin{aligned} \text{minimize} \quad & p_1 + p_2 + p_3 \\ \text{subject to} \quad & p_1 \geq 9 \\ & p_2 \geq 5.5 \\ & p_3 \geq 5.5 \\ & p_1 + p_2 \geq 15 \\ & p_1 + p_3 \geq 15 \\ & p_2 + p_3 \geq 11 \\ & p_1 + p_2 + p_3 \geq 15 \end{aligned}$$

Recall that each of the constraints is defined in Equation (3) of Section 2. For the instance above, it is easy to compute the minimum revenue of the auctioneer (the value of the objective funtion) by taking all the combinations of the MRCS constraints and observing that one set of constraints that must be satisfied by equality are $p_2 \geq 5.5$ and $p_1 + p_3 \geq 15$ which, combined yield that

$$(p_1 + p_3) + (p_2) \geq 20.5.$$

Hence, $\mathsf{MREV}(\mathbf{b}) = 20.5$.

Now suppose that bidder 1 unilaterally declares $b'_i(\{A\}) = 10 > 9.5$. The optimal allocation remains $X^*(b_1, \mathbf{p}_{-1}) = \{1, 2, 3\}$. However, the new MRCS LP becomes:

$$\begin{aligned} \text{minimize} \quad & p_1 + p_2 + p_3 \\ \text{subject to} \quad & p_1 \geq 9 \\ & p_2 \geq 5 \\ & p_3 \geq 5 \\ & p_1 + p_2 \geq 15 \\ & p_1 + p_3 \geq 15 \\ & p_2 + p_3 \geq 11 \\ & p_1 + p_2 + p_3 \geq 15 \end{aligned}$$

Once again, we consider all combinations of constraints. we see that since this time the VCG constraints of bidders 2 and 3 have been *relaxed* ,the "blocking" constraints are all greater than all or equal to 20 or strictly weaker. Hence, $\mathsf{MREV}(b'_i, \mathbf{b}_{-i}) = 20 < 20.5 = \mathsf{MREV}(\mathbf{b})$ and the proof follows. Note that this example also implies that the core polytope strictly increases its solution space (Theorem 1). For example, the solution $10, 5, 5$ is now feasible, whereas when bidder 1 was bidding 9.5 this was not possible as it violated individual rationality ☐

B Missing Proofs from Section 4

Proof of Lemma 9.

Given a profile $\mathbf{v}$, fix a bidder $i \in \hat{X}(\mathbf{v})$. Suppose this bidder unilaterally declares a bid $v'_i > v_i$. We will show that bidder i remains in the set of final winners $\hat{X}(v'_i, \mathbf{v}_{-i})$ for $v'_i > v_i$. Recall that the mechanism we propose, given a profile $\mathbf{v}$ finds an initial provisional allocation $X^*(\mathbf{v})$ and then selects a $\hat{X}(\mathbf{v}) \subseteq X^*(\mathbf{v})$. Hence, for all $v'_i > v_i$ we need to argue both that $i \in X^*(v'_i, \mathbf{v}_{-i})$ and $i \in \hat{X}(v'_i, \mathbf{v}_{-i})$.

For the first step, monotonicity is implied by the fact that the allocation algorithm $\hat{X}$ calls X^*, and since the welfare-maximizing algorithm X^* is monotone it holds that $i \in X^*(v'_i, \mathbf{v}_{-i})$ for all $v'_i > v_i$.

For the second step, Since bidder i has unilaterally declared a bid $v'_i > v_i$ we can be certain her index in the ranking among bidders in $X^*(\mathbf{v}) = X^*(v'_i, \mathbf{v}_{-i})$ can only be lower. By Lemma 8 we know that $|\hat{X}(\mathbf{v})| \leq |\hat{X}(v'_i, \mathbf{v}_{-i})|$, which implies that the allocation algorithm has picked a superset of $\hat{X}(\mathbf{v})$. Nevertheless, bidder i will be a part of the new optimal allocation $\hat{X}(v'_i, \mathbf{v}_{-i})$. ☐

Proof of Lemma 11.

Given a vector $\mathbf{v}$, the total revenue of the auctioneer is the sum of the threshold payments of bidders in $\hat{X}(\mathbf{v})$. Recall that for every bidder $i \notin \hat{X}(\mathbf{v})$, by Myerson's Lemma $\hat{p}_i(\mathbf{v}) = 0$. Hence, we can lower bound the total revenue of the auctioneer as follows:

$$\sum_{j\in\hat{X}(\mathbf{v})} \hat{p}_j(\mathbf{v}) \geq \sum_{i\in\hat{X}(\mathbf{v})} \frac{\mathsf{MREV}(\hat{p}_i(\mathbf{v}), \mathbf{v}_{-i})}{|\hat{X}(\hat{p}_i(\mathbf{v}), \mathbf{v}_{-i})| \cdot H_n} \geq \sum_{i\in\hat{X}(\mathbf{v})} \frac{\mathsf{MREV}(\mathbf{v})}{|\hat{X}(\hat{p}_i(\mathbf{v}), \mathbf{v}_{-i})| \cdot H_n}$$
$$\geq \sum_{i\in\hat{X}(\mathbf{v})} \frac{\mathsf{MREV}(\mathbf{v})}{|\hat{X}(\mathbf{v})| \cdot H_n} = \frac{\mathsf{MREV}(\mathbf{v})}{H_n}.$$

The first inequality follows from Lemma 10 (Equation (14)). To obtain the second inequality, for every bidder $i \in \hat{X}(\mathbf{v})$ we apply Corollary 2 for the profile $(\hat{p}_i(\mathbf{v}), \mathbf{v}_{-i})$. Note that the vector $(\hat{p}_i(\mathbf{v}), \mathbf{v}_{-i})$ satisfies the conditions of Corollary 2 since, by Lemma 10 it is also true that $\hat{p}_i(\mathbf{v}) \geq p_i^{VCG}(\mathbf{v})$. Finally, we obtain the third inequality by applying Lemma 8 regarding the same profile and deviation. ◻

C Missing Proofs from Section 5

Proof of Theorem 4. For a profile $\mathbf{b}$, let $\mathbf{p}(\mathbf{b})$ be the payment vector of the $\mathbf{f}(\mathbf{b})$-nearest rule. Suppose for a contradiction that (17) is not satisfied, i.e., there exists a profile $\mathbf{b}$, a bidder $i \in X^*(\mathbf{b})$, and a bid $b'_i > b_i$ for which

$$p_i(b'_i, \mathbf{b}_{-i}) < p_i(\mathbf{b}). \tag{20}$$

Note that due to the monotonicity of X^*, it holds that $X^*(\mathbf{b}) = X^*(b'_i, \mathbf{b}_{-i})$ for $b'_i > b_i$. Since $\mathbf{f}(\mathbf{b})$-nearest is a Pareto-efficient payment rule, by Lemma 3, for the deviating bidder i it holds that $p_i(\mathbf{b}) \leq b_i$. By combining this fact along with (20), we obtain that $p_i(b'_i, \mathbf{b}_{-i}) < b_i$. In its turn, by Theorem 2 we have that

$$\mathbf{p}(b'_i, \mathbf{b}_{-i}) \in CORE(\mathbf{b}). \tag{21}$$

Equation (21) implies that the optimal solution for the profile $(b'_i, \mathbf{b}-i)$ is actually a member of the initial core polyhedron defined for the vector $\mathbf{b}$. We distinguish that each of the following cases leads to a contradiction:

1. $p_i(\mathbf{b}) = p_i^{VCG}(\mathbf{b})$: Bidder i cannot be asked to pay a payment $p_i(b'_i, \mathbf{b}_{-i}) < p_i^{VCG}(\mathbf{b})$ as this would violate the constraint for coalition $N \setminus \{i\}$ in (3). This would imply that $\mathbf{p}(b'_i, \mathbf{b}_{-i}) \notin CORE(b'_i, \mathbf{b}_{-i})$.
2. $p_i(\mathbf{b}) > p_i^{VCG}(\mathbf{b})$ *and* $\mathsf{MREV}(b'_i, \mathbf{b}_{-i}) < \mathsf{MREV}(\mathbf{b})$: In this case, the solution $\mathbf{p}(b'_i, \mathbf{b}_{-i})$ achieves a strictly lower minimum revenue when compared to $\mathbf{p}(\mathbf{b})$. However, by (21), $\mathbf{p}(b'_i, \mathbf{b}_{-i}) \in CORE(\mathbf{b})$, which is a contradiction since it implies that the solution $\mathbf{p}(\mathbf{b})$ is not an MRCS solution of $CORE(\mathbf{b})$.
3. $p_i(\mathbf{b}) > p_i^{VCG}(\mathbf{b})$ *and* $\mathsf{MREV}(b'_i, \mathbf{b}_{-i}) = \mathsf{MREV}(\mathbf{b})$: To analyze this case, let us first define for every vector $\mathbf{r}$, the function

   $$D(\mathbf{b}, \mathbf{r}) := \sum_{j\in X^*(\mathbf{b})} (p_j(\mathbf{b}) - r_j)^2.$$

   For a given profile $\mathbf{b}$, $D(\mathbf{b}, \mathbf{f}(\mathbf{b}))$ is the optimal value of the objective function of $\mathbf{f}(\mathbf{b})$-nearest. We have now the following implications:

   $$\begin{aligned} D((b'_i, \mathbf{b}_{-i}), \mathbf{f}(b'_i, \mathbf{b}_{-i})) &< D(\mathbf{b}, \mathbf{f}(b'_i, \mathbf{b}_{-i})) \\ &= D(\mathbf{b}, \mathbf{f}(\mathbf{b})) + (f_i(b'_i) - f_i(b_i))\,(f_i(b'_i) + f_i(b_i) - 2p_i(\mathbf{b})) \\ &< D((b'_i, \mathbf{b}_{-i}), \mathbf{f}(\mathbf{b})) + (f_i(b'_i) - f_i(b_i))\,(f_i(b'_i) + f_i(b_i) - 2p_i(\mathbf{b}))\,. \end{aligned} \tag{22}$$

   The first inequality follows from the fact that since $\mathbf{p}(b'_i, \mathbf{b}_{-i})$ is the unique optimal solution for $CORE(b'_i, \mathbf{b}_{-i})$, and since, by Theorem 1, $\mathbf{p}(\mathbf{b})$ is also a feasible payment in $CORE(b'_i, \mathbf{b}_{-i})$, the value of the objective function $D(\mathbf{b}, \mathbf{f}(b'_i, \mathbf{b}_{-i}))$ must be strictly larger. We apply the same argument and obtain the last inequality (Equation (22)) for the $CORE(\mathbf{b})$ polyhedron, since by Equation (21), $\mathbf{p}(b'_i, \mathbf{b}_{-i}) \in CORE(\mathbf{b})$. By rearranging terms, Equation (22) implies

   $$(f_i(b'_i) - f_i(b_i))\,(2p_i(\mathbf{b}) - 2p_i(b'_i, \mathbf{b}_{-i})) < 0, \tag{23}$$

   a contradiction, since for $b'_i > b_i$ we have that $f_i(b'_i) \geq f_i(b_i)$ by the monotonicity of $f_i(\cdot)$ and, by assumption $p_i(\mathbf{b}) > p_i(b'_i, \mathbf{b}_{-i})$.

□

Detailed Derivation of Equation (23) **from Equation** (22)

$$
\begin{aligned}
&D((b_i', \mathbf{b}_{-i}), \mathbf{f}(b_i', \mathbf{b}_{-i})) < D((b_i', \mathbf{b}_{-i}), \mathbf{f}(\mathbf{b})) + (f_i(b_i') - f_i(b_i))\,(f_i(b_i') + f_i(b_i) - 2p_i(\mathbf{b})) \Leftrightarrow \\
&((p_i(b_i', \mathbf{b}_{-i})) - f(b_i'))^2 < ((p_i(b_i', \mathbf{b}_{-i})) - f(b_i))^2 + (f_i(b_i') - f_i(b_i))\,(f_i(b_i') + f_i(b_i) - 2p_i(\mathbf{b})) \Leftrightarrow \\
&((p_i(b_i', \mathbf{b}_{-i})) - f(b_i'))^2 - ((p_i(b_i', \mathbf{b}_{-i})) - f(b_i))^2 < (f_i(b_i') - f_i(b_i))\,(f_i(b_i') + f_i(b_i) - 2p_i(\mathbf{b})) \Leftrightarrow \\
&(f(b_i') - f(b_i))(f(b_i') + f(b_i) - 2p_i(b_i', \mathbf{b}_{-i})) < (f_i(b_i') - f_i(b_i))\,(f_i(b_i') + f_i(b_i) - 2p_i(\mathbf{b})) \Leftrightarrow \\
&(f_i(b_i') - f_i(b_i))\,(2p_i(\mathbf{b}) - 2p_i(b_i', \mathbf{b}_{-i})) < 0.
\end{aligned}
$$