# Obviously Strategyproof Single-Minded Combinatorial Auctions


**Bart de Keijzer**
King's College London, UK
bart.de_keijzer@kcl.ac.uk

**Maria Kyropoulou**
University of Essex, UK
maria.kyropoulou@essex.ac.uk

**Carmine Ventre**
King's College London, UK
carmine.ventre@kcl.ac.uk



## Abstract

We consider the setting of combinatorial auctions when the agents are single-minded and have no contingent reasoning skills. We are interested in mechanisms that provide the right incentives to these imperfectly rational agents, and therefore focus our attention to *obviously strategyproof (OSP)* mechanisms. These mechanisms require that at each point during the execution where an agent is queried to communicate information, it should be "obvious" for the agent what strategy to adopt in order to maximise her utility. In this paper we study the potential of OSP mechanisms with respect to the approximability of the optimal social welfare.

We consider two cases depending on whether the desired bundles of the agents are *known* or *unknown* to the mechanism. For the case of known-bundle single-minded agents we show that OSP can actually be as powerful as (plain) strategyproofness (SP). In particular, we show that we can implement the very same algorithm used for SP to achieve a $\sqrt{m}$-approximation of the optimal social welfare with an OSP mechanism, $m$ being the total number of items. Restricting our attention to declaration domains with two values, we provide a 2-approximate OSP mechanism, and prove that this approximation bound is tight. We also present a randomised mechanism that is universally OSP and achieves a finite approximation of the optimal social welfare for the case of arbitrary size finite domains. This mechanism also provides a bounded approximation ratio when the valuations lie in a bounded interval (even if the declaration domain is infinitely large). For the case of unknown-bundle single-minded agents, we show how we can achieve an approximation ratio equal to the size of the largest desired set, in an OSP way. We remark this is the first known application of OSP to multi-dimensional settings, i.e., settings where agents have to declare more than one parameter.

Our results paint a rather positive picture regarding the power of OSP mechanisms in this context, particularly for known-bundle single-minded agents. All our results are constructive, and even though some known strategyproof algorithms are used, implementing them in an OSP way is a non-trivial task.



**2012 ACM Subject Classification** Theory of computation → Algorithmic game theory and mechanism design

**Keywords and phrases** OSP Mechanisms, Extensive-form Mechanisms, Single-minded Combinatorial Auctions, Greedy algorithms

**Digital Object Identifier** 10.4230/LIPIcs.ICALP.2020.71

**Category** Track A: Algorithms, Complexity and Games

# 1 Introduction

*Algorithms might not have direct access to their inputs.* This is by now a well-known issue, motivated by the internet and the self interests thereon. A body of work in computer science has been devoted to the design of algorithms that can faithfully elicit the input from their selfish sources. Typically, this is achieved through so-called mechanism design, a subdiscipline of game theory that studies the design of functions, or *mechanisms*, where the input is provided by multiple self-interested agents, and the output, or *outcome*, needs to satisfy a set of pre-specified desirable objectives. Each agent has a certain utility for each outcome. In its most general definition, a mechanism is given by a multi-round interaction protocol between a central authority and the agents, defining a game wherein the central authority wants to compute a certain function of the agents' inputs and the agents choose a strategy leading to the outcome that maximises their own utility. The main design requirements for such mechanisms are:

**Strategyproofness (SP).** There is a strategy for each agent that is guaranteed to result in a better outcome than any other strategy that the agent may adopt, irrespective of the other agents' strategies.

**(Approximate) economic efficiency.** The outcome of the mechanism must have a quality that is close to a theoretical optimum; this is often measured in terms of *social welfare*, i.e., the sum of all the agents' utilities.

Combinatorial auctions (CAs) have emerged as the paradigmatic problem in the area, exemplifying the tension between these two desiderata and the polynomial-time computation of the outcome. In CAs, we are given a set of items that need to be sold among a set of agents who are interested in buying the items. These agents express valuations for each of the items, or certain bundles of items, or even each possible bundle of items, depending on the level of complexity of the utility model that is assumed. A mechanism in this setting must determine how the items are allocated to the agents, and how much each agent is going to be charged.

Whilst it is not known to what extent it is possible to guarantee the properties above with polynomial-time algorithms, the design of mechanisms that satisfy these objectives is reasonably well-understood for some sufficiently simple auction settings. This is the case for the optimization problem of interest to this study, *single-minded combinatorial auctions*. In such a setting, there are multiple agents and multiple items to be allocated. An agent's utility function has a simple form: It is determined by a *valuation* of the set of items that the agent gets allocated, from which a potential payment charged by the mechanism is subtracted. The valuation function has the following structure: For an agent, say $i$, there is a number $v_i$ and a subset of items $R_i$ such that her valuation is $v_i$ whenever the allocated bundle contains $R_i$, and 0 otherwise. The utility of each agent is therefore determined by a single pair $(v_i, R_i)$. This simple setting has been the central object of study of the celebrated paper [17], where the authors design a simple greedy mechanism that is strategyproof and approximates the optimal social welfare to within a factor of $\sqrt{m}$, where $m$ is the number of items. It is well-known that this combinatorial optimisation problem is NP-hard and inapproximable within a factor of $\sqrt{m}$ unless $\mathsf{NP} = \mathsf{ZPP}$ [14].

In many important cases, however, the implementation of strategyproof mechanisms can be too complex, unintelligible, unintuitive, or cognitively too demanding due to the limited ability of typical agents, such as, the capability to carry out *contingent reasoning*. Even for the simple setting of one-item auctions, a special case of single-minded agents, it is known how implementation details matter, see, for example, [2], for a discussion on the differences between sealed bid versus ascending bid auctions, and [15] for a related experimental study. We refer furthermore to [6] and [4] for further reading on the issue of contingent reasoning. Hence, the theoretically strong mechanisms that have been proposed in past mechanism design literature may be too difficult to understand and to use for agents with imperfect rationality.

This problem motivates the design of mechanisms with a reduced *cognitive burden* for agents who participate in the mechanism: Mechanisms should be very easy to understand, and transparent to participate in. While it is somewhat of a challenge to satisfactorily define formally what a "low cognitive burden" comprises, the concept of *obvious strategyproofness* (OSP) which has been introduced in [18] provides a strong and reasonably satisfactory notion in the context of agents unable to reason contingently. Informally, OSP requires that at each point during the execution of a mechanism where an agent is queried to reveal information (i.e., point where an agent is asked to make a decision about how the execution of the mechanism should proceed), it should be "obvious" for the agent what strategy to adopt in order to maximise her utility: For the agent, there must be a single choice for which it holds that all outcomes that can result from that choice are better for her than any other outcome that can result from an alternative choice. OSP therefore strengthens the classical notion of strategyproofness.

**Our contributions.** We study the extent to which OSP mechanisms can return good approximations to the optimal social welfare in the setting of single-minded combinatorial auctions. We measure the quality of mechanisms in terms of a relative approximation ratio. As in much of the literature on OSP, we assume that the set of possible types that the agents can have (which we refer to as the *declaration domain*) is publicly known; this is either finite or contained in a closed interval. As for the agents, we prove results depending on whether they are *known-bundle single-minded* – whereby the valuation $v_i$ for agent $i$'s desired bundle $R_i$ is not known but $R_i$ is – or, *unknown-bundle single-minded*, for which neither $v_i$ nor $R_i$ is known and must be elicited in an OSP way.

For the case of known-bundle single-minded agents, we provide the following results.

- If there are only two possible valuations, i.e., a low valuation $L$ and a high valuation $H$, we express the characterisation of OSP mechanisms in [7] conveniently and design a deterministic OSP mechanism with an approximation ratio of 2, which we show to be the best possible. We give explicit payment functions for this mechanism and we can prove that truthtelling agents always have non-negative utility (a property known as individual rationality (IR)). We furthermore show that if the OSP mechanism were to use a fixed ordering of agents that does not depend on the instance, then the approximation guarantee of the mechanism is unbounded.
- If the declaration domain is an arbitrarily large finite set, we derive an OSP mechanism that achieves an approximation ratio of $\sqrt{m}$ to the optimal social welfare. This mechanism can be regarded as an obviously strategy-proof implementation of the mechanism in [17]. The payments here are implicitly given through the *cycle monotonicity* technique developed in [7]. Our mechanism makes use of the fact that the domain is finite and that the desired bundles are known.
- We further provide a randomised OSP and IR mechanism that achieves an approximation ratio strictly less than $d$, where $d$ is the cardinality of the declaration domain. In particular, for arbitrary size finite domains of $d$ valuations $V_1 < \cdots < V_d$, our mechanism achieves a $(d - V_1/V_2 - \cdots - V_{d-1}/V_d)$-approximation of the optimal social welfare. The idea behind this mechanism is to simply draw at random (with a carefully chosen probability

distribution) one value from the domain and ask agents if their valuation is at least as big; the mechanism thus stays OSP even if agents had access to the randomness used (i.e., they are "universally OSP"). We note that this result improves the aforementioned bound of 2 for two-value domains, and approaches 2 as $V_1/V_2 \to 0$. We also generalise this result for (uncountable) valuations contained in an interval $[a, b]$ and derive the continuous limit of the above mechanism, which results in an OSP and IR mechanism for this setting that achieves an approximation ratio of $1 + \ln(b/a)$ to the optimal social welfare. This means that the approximation ratio only grows logarithmically in the relative size of the interval.

The results above paint a rather positive picture regarding the power of OSP mechanisms for CAs with known-bundle single-minded agents. Firstly, our upper bound of $\sqrt{m}$, for finite valuation domains, matches what is known for strategyproofness – this is, to best of our knowledge, the first case in which OSP has been proved to be *as powerful as* SP. ([16] proves an asymptotic equivalence for randomised OSP without money for a variant of a scheduling problem.) The only additional time we need is used to sweep through the declaration domain of the agents, a seemingly unavoidable step to guarantee an OSP implementation of direct-revelation mechanisms. Secondly, this result shows that Deferred Acceptance (DA) auctions, whilst being OSP, are not the right algorithmic approach to optimise the approximation guarantee of OSP mechanisms. We in fact beat the lower bounds proved in [5] regarding the approximation guarantee of DA auctions in this setting. This observation reinforces the findings in [7, 9] concerning the power of DAs vs OSP, in the context of scheduling related machines. Thirdly, our tight bounds for two-value domains show (i) how the graph-theoretic approach to OSP in [7] can be made operational; and, (ii) that the order in which the mechanism queries the set of agents is of crucial importance in the design of the mechanism. Finally, our randomised OSP mechanisms show how it is possible to leverage "internal" chance nodes [18] and beat deterministic mechanisms, whilst agents do not need to compute expectations. This is to our knowledge the first known application of randomisation over OSP mechanisms.

OSP has so far only been considered for single-parameter settings, i.e., where the agents' private information is a single number, as no explicit technique is known to implement OSP mechanisms for higher dimensional settings. However, we complete this paper by giving the *first* OSP mechanism for agents with richer declaration domains. We in fact provide an OSP mechanism for the case of unknown-bundle single-minded agents, that returns an approximation of the optimal social welfare equal to the maximum size of a desired set $R_i$. This is an OSP implementation of the well-known Greedy-by-valuation algorithm for CAs. To obtain this result, we leverage the cycle monotonicity characterisation of OSP in [7] and an approach recently used in [10] to deal with long negative cycles. The idea is to give a structural property of the negative-weight cycles that have more than two vertices, for mechanisms that satisfy a natural notion of monotonicity between two instances. We then prove that any such mechanism that additionally queries the agents monotonically (that is, from an extreme of the domain to the other, irrespective of the desired set) cannot have any long negative-weight cycle. This is only proved to be a sufficient condition; the extent to which this is also necessary (proved to be true for single-parameter settings in [10]) will say whether our bound can be improved or not. This is the main open problem left by our work.

**Further Related Work.** The notion of OSP introduced in [18] has spawned numerous subsequent studies in both the computer science and economics community.

In [1], the authors study the OSP concept in the context of stable matchings and provide a suitable mechanism under an acyclicity assumption, as well as an impossibility result for a more general setting. The paper [3] further studies this concept for housing markets, single-peaked domains, and a general quasi-linear mechanism design setting. In [20] the authors study OSP mechanisms in general design domains where monetary transfers are not possible, and they provide a useful characterisation of OSP mechanisms in this setting.

OSP was investigated in machine scheduling domains and set system problems in [7], published in [9, 8], and furthermore a study was done in [11] where the authors prove that a restriction on the agents' declaration called *monitoring* can help obtain OSP mechanisms with a good approximation ratio in various mechanism design domains. The paper [16] builds forth on this by studying machine scheduling in the absence of monetary payments. For machine scheduling, a generalisation of OSP is furthermore studied in [13] where the restriction on the ability of agents to reason contingently is weakened, and the authors show that a large amount of "look-ahead"-ability is required for the agents in any mechanism that achieves a good approximation ratio in the considered scheduling setting. Another study that considers OSP under a restriction on the agents' behavior is [12], where the authors assume that non-truthful behaviour can be detected and penalised with a certain probability.

In [19] a revelation principle is presented that states that every social choice function implementable through an OSP mechanism can be implemented using a certain structured OSP protocol where agents take turns making announcements about their valuations.

## 2 Preliminaries

In a combinatorial auction we have a set $U$ of $m$ items and a set $N$ of $n$ agents. Each agent $i$ has a *private* valuation function $v_i$ and, in the general case, is interested in obtaining only one set in a *private* collection $\mathcal{S}_i$ of subsets of $U$, also called bundles. Thus, the valuation function maps subsets of items to nonnegative real numbers ($v_i(\emptyset)$ is normalised to be 0). The agents' valuations are monotone: for $S \supseteq T$ we have $v_i(S) \geq v_i(T)$. In single minded combinatorial auctions, $|\mathcal{S}_i| = 1$, each agent $i$ is interested in obtaining only one particular subset of $U$; we denote $i$'s desired bundle by $R_i$. This implies that agent $i$'s valuation is the same for all supersets of $R_i$, while it is 0 otherwise. Formally, consider agent $i$ and let $\mathcal{S}_i = \{R_i\}$. The valuation function of agent $i$ for a given set $S$ is

$$v_i(S) = \begin{cases} v_i & \text{if } S \supseteq R_i, \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

where (with a slight notation overloading) $v_i \in \mathbf{R}_{\geq 0}$ is a non-negative real number. Note that the valuation function of an agent $i$ is fully represented by her desired set $R_i$ and her valuation for that set $v_i$. The goal is to find a partition $(S_1, \ldots, S_n)$ of $U$ such that $\sum_{i=1}^{n} v_i(S_i)$ –the *social welfare (SW)*– is maximised. We denote the optimal (maximum possible) social welfare by $SW^*$ (we omit the dependence on the instance as it will be clear from the context).

We consider two versions of single-minded combinatorial auctions. In the case of *unknown-bundle single-minded agents*, we assume that the desired sets $R_i$ and the valuations $v_i$ are private knowledge of the agents, while in the *known-bundle single-minded agents* case, we assume that the desired sets $R_i$ are known and only the valuations $v_i$ are private knowledge of the agents. In each case, we refer to the private information of an agent as her *type*. Agents are typically asked to declare their types. We use $D_i$, the *declaration domain* of agent $i$, to denote the set of all possible types of agent $i$. In this paper we will consider declaration domains that are identical for all agents $i$, i.e., $D_i = D$. We use $\mathbf{b} = (b_1, \ldots, b_n)$ to denote declaration profiles, so that $b_i \in D$ stands for an agent $i$'s demanded set $R_i$ along with the valuation $v_i$ that she has for it.

We want to design an auction mechanism that interacts with the agents in any sequence, and outputs (based on this interaction) an allocation $A_i$ of items to each agent $i$, along with a payment $P_i > 0$ that is charged to each agent $i$. The *allocation function* satisfies that each item is allocated to exactly one agent. We refer to such a pair $(A = (A_1, \dots, A_n), P = (P_1, \dots P_n))$ as an *outcome*. We assume that agents have quasi-linear utility functions, that is, agent $i$ with type $b_i$ has a utility of

$$u_i(b_i, \mathbf{b}_{-i}) = v_i(A_i) - P_i$$

for outcome $(A, P)$. Each agent will interact strategically with the mechanism, so as to make it output an outcome that maximises her utility.

When designing the mechanism, we want to define allocations and payments in such a way that they induce a certain type of behaviour of the strategically acting agents, while simultaneously ensuring that the outcome maximises or approximately maximises the social welfare. In particular, we aim to design *obviously strategyproof (OSP)* mechanisms. To define this notion, we view a mechanism in the form of an *implementation tree* $\mathcal{T}$ that captures the way in which the mechanism interacts with the agents [18, 7].

We now introduce some notation around $\mathcal{T}$ before we formally define the OSP notion. Each internal node $u$ of $\mathcal{T}$ is labeled with an agent $\mathcal{Q}(u)$, called the *divergent agent* at $u$, and the outgoing edges from $u$ are labeled with sets of types in the declaration domain of $\mathcal{Q}(u)$. At node $u$, the agent $\mathcal{Q}(u)$ is queried and asked to choose an action, that corresponds to selecting one of $u$'s outgoing edges. The labels of the outgoing edges of $u$ form a partition of the *current domain* of $i$, denoted as $D_i(u)$. The current domain $D_i(u)$ is equal to the label of the last edge $e$ in the path from the root to $u$ that $i$ chose as an action. When player $\mathcal{Q}(u)$ chooses an outgoing edge at node $u$, we say that the chosen action *signals* that the type of $\mathcal{Q}(u)$ is in the set of types labeling the chosen edge. If a pair of types in the current domain of agent $i$ occur in the labels of two distinct outgoing edges of $u$, we say that $i$ is asked to *separate* the two possible types at $u$. The leaves of the tree will thus be linked to (a set of) type profiles; and at each leaf the mechanism will return an outcome $(A, P)$ accordingly; in other words, each leaf corresponds to an outcome of the mechanism. (Observe that this means that the domain of $A$ and $P$ is effectively given by the leaves of $\mathcal{T}$.)

A mechanism (viewed in the form of an implementation tree as described above) is said to be OSP if at each node where an agent is asked to diverge, she *always* maximises her utility by choosing the edge $(u, u')$ with the label containing her own type, in every node. That is, if we call the latter strategy $s$, then a mechanism is said to be OSP if the worst possible outcome after signaling her true type (taken over all the reachable outcomes in the leaves of the subtree rooted at $u'$, with respect to $s$) gives her a utility *at least as good* as when she would get the best possible outcome after choosing any other edge $(u, u'')$ at that particular point in the implementation tree (taken over all the outcomes in the leaves of the subtree rooted at $u''$). The corresponding utility-maximising strategies played by the players are called the *OSP strategies*. In cases where we discuss any particular OSP mechanism, we use $(A(\mathbf{b}), P(\mathbf{b}))$ to denote the outcome (i.e., allocation and payment vector) that results from agents playing their OSP strategies, i.e., strategy profiles where each agent at every node in the mechanism follows the edge containing her type.

We call a type profile $\mathbf{b}$ *compatible with* $u \in \mathcal{T}$ iff for each edge $(u', u'')$ on the path from the root to $u$, the label of $(u', u'')$ contains $b_{\mathcal{Q}(u')}$. We furthermore say that two type profiles $(t, \mathbf{b}_{-i})$ and $(b, \mathbf{b}'_{-i})$ *diverge* at $u$ if $i = \mathcal{Q}(u)$ and $t$ and $b$ are labels of different edges outgoing from $u$ (we sometimes will abuse this terminology and we also say that $t$ and $b$ diverge at $u$). For every agent $i$ and types $t, b \in D_i$, we let $u^i_{t,b}$ denote a vertex $u$ in the implementation tree $\mathcal{T}$, such that $(t, \mathbf{b}_{-i})$ and $(b, \mathbf{b}'_{-i})$ are compatible with $u$, but diverge at $u$ for some $\mathbf{b}_{-i}, \mathbf{b}'_{-i} \in D_{-i}(u) = \times_{j \neq i} D_j(u)$. Note that such a vertex might not be unique as agent $i$ will be asked to separate $t$ from $b$ in different paths from the root (but only once for every such path). We call these vertices of $\mathcal{T}$ *tb-separating* for agent $i$. For example, the node $r$ in the tree in Figure 1 is an $LH$-separating node for agent 1; while $v$ and $w$ are two $LH$-separating nodes for agent 2.

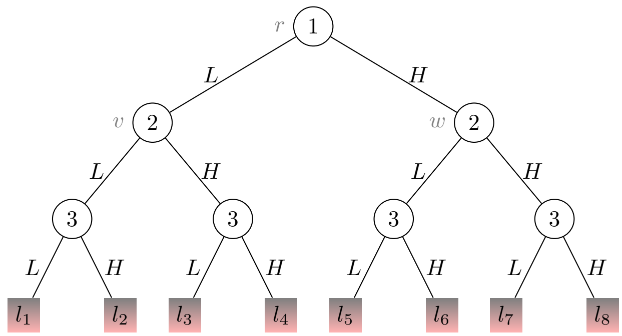


**Figure 1** An implementation tree with three agents with two-value domains $\{L, H\}$; each agent separates the domain types upon playing; at each leaf $l_i$ the mechanism computes $A(\mathbf{b})$ and $P(\mathbf{b})$, $\mathbf{b}$ being the declaration vector at $l_i$.

We also consider randomised mechanisms that are *universally OSP*, in Section 5. Such mechanisms are probability distributions on OSP mechanisms.

Besides the OSP requirement, we would like our mechanisms to satisfy *(ex-post) individual rationality (IR)*, i.e., when an agent plays an OSP strategy, the agent is guaranteed an outcome that gives her a non-negative utility. We measure the quality of mechanisms in terms of a relative *approximation ratio*, i.e., an upper bound on the ratio of the optimal social welfare and the social welfare of the outcome of the mechanism.

# 3 Deterministic mechanisms for known-bundle single-minded agents

We now focus on the case of known-bundle single-minded agents, i.e. the desired sets of the agents are known by the mechanism, and only the valuations of the agents for their corresponding desired set is private information. We let $v_i, i \in N$ denote the valuation of agent $i$ for her desired bundle, and slightly abusing notation, we use $SW(I) = \sum_{i \in I} v_i$ where $I$ is a set of agents who can be allocated their desired bundles at the same time.

## 3.1 Domains of size 2

We first restrict attention to the case where the declaration domain of each agent has two possible values and begin with a simple and elegant characterisation of individually rational OSP mechanisms for the case of known-bundle single-minded agents with two-value domains $D_i = \{L, H\}$, for $i \in N$.

▶ **Definition 1** ($\mathcal{C}$ mechanisms). *We define a class $\mathcal{C}$ of mechanisms, that only use the following types of queries:*

- $L$-query: *The divergent agent is asked to separate $L$ and $H$. If she signals $L$ then she will not get her desired bundle (regardless of the future).*
- $H$-query: *The divergent agent is asked to separate $L$ and $H$. If she signals $H$ then she is guaranteed to get her desired bundle.*

▶ **Theorem 2.** *The class $\mathcal{C}$ of mechanisms characterises deterministic IR and OSP mechanisms for the case where the declaration domain is $\{L, H\}$ and the desired bundles are known, in the following sense:*

- *There exist payments such that every mechanism in $\mathcal{C}$ is deterministic IR and OSP.*
- *Every deterministic IR OSP mechanism is equivalent to a mechanism in $\mathcal{C}$ with respect to their allocations, i.e. for every possible valuation profile of the agents, the allocations resulting from both mechanisms are identical.*

**Proof.** Regarding the first claim, fix the payment of an agent who is allocated her bundle to be $L$, and the payment to be 0 otherwise. It should be straightforward to see why pairing mechanisms in $\mathcal{C}$ with these payment yields a deterministic IR mechanism. Moreover, observe that regardless of the query, if an agent has valuation $L$, then her utility will be 0 regardless of her signals and allocation (even if she is allocated her bundle, then she will be asked to pay $L$). If an agent has valuation $H$ and she is asked an $L$-query, she can only get positive utility for signaling her true valuation $H$, while if she is asked an $H$-query, she can guarantee herself the maximum possible utility under the mechanism $(H - L)$ again by reporting her true valuation $H$. This demonstrates that whenever an agent is asked to diverge, she is not worse off by acting according to her true type in *any* possible future scenario, hence OSP is satisfied.

Regarding the second claim, consider any OSP mechanism and its associated implementation tree $\mathcal{T}$. Clearly, since we are dealing with the case of two-type domains, if $\mathcal{T}$ has a node with more than two children, then this node can be pruned to yield an OSP mechanism where the node has exactly two children. Furthermore, nodes that have only one child are trivial and can be removed from the tree. So, at every node in $\mathcal{T}$ an agent is asked to separate $L$ and $H$, and at any path from the root to a leaf in $\mathcal{T}$ each agent is only asked to diverge at most once.

Next, fix a node $u \in \mathcal{T}$ and let agent $i = \mathcal{Q}(u)$ be the divergent agent at $u$. Suppose for a contradiction that if $i$ signals $L$ when queried at $u$, then it is possible that $i$ gets her desired bundle, while if $i$ signals $H$ then it is possible that $i$ is not allocated her desired bundle. In other words, the corresponding subtrees of $\mathcal{T}$ have outcomes $o_L$ where $i$ gets her desired bundle when signaling $L$, and $o_H$ where $i$ doesn't get her desired bundle when signaling $H$. In this case the OSP condition would be violated at node $u$ when $v_i = H$, as $i$'s utility under $o_H$ is lower than under $o_L$ (by IR the payment under $o_H$ is 0, while the payment under $o_L$ is at most $L$). We can conclude that at any node $w \in \mathcal{T}$, either all possible outcomes when the divergent agent $j = \mathcal{Q}(w)$ signals $L$ do not allocate $j$'s desired bundle to $j$, or all outcomes when $j$ signals $H$ allocate $j$'s desired bundle to $j$. These two alternatives correspond to the $L$-queries and $H$-queries in the definition of $\mathcal{Q}$. ◀

Next, we use the above characterisation theorem[1] to provide a tight bound on the approximability of the optimal social welfare under OSP mechanisms.

▶ **Theorem 3.** *The $IS$ mechanism described in Algorithm 1 is an OSP mechanism that achieves an approximation ratio $\rho \leq 2$ for the case of known-bundle single-minded agents and domains of size* 2.

**Proof.** Consider any instance of the problem, let $I \subseteq N$ be the allocation output by $IS$, and let $I^* \subseteq [n]$ be the SW-maximising allocation according to the true valuations. Also fix $\mathcal{N}$ to its value at the last iteration of IS. At the last iteration of $IS$ it holds that every agent in

[1] Note that Theorem 2 essentially expresses the cycle monotonicity property of [7] in a convenient way.

**Algorithm 1** The $IS$ mechanism.

1 $\mathcal{N} \leftarrow N$ ($\mathcal{N}$ is the set of agents currently under consideration) $A_q \leftarrow \emptyset$ ($A_q$ is the set of agents who have already been queried) $I \leftarrow \emptyset$ ($I$ is the set of agents currently allocated their desired bundle) $v_i = L$, for $i \in N$ (Our mechanism originally assumes that all agents have valuation $L$)
2 Compute the SW-maximising feasible allocation and update set $I$ accordingly
3 **while** *there exists an agent* $i \in \mathcal{N} \setminus A_q$ *who is not in* $I$ **do**
4     Perform an $L$-query to $i$ (break ties arbitrarily)
5     **if** $i$ *signals* $L$ **then**
6         $\mathcal{N} = \mathcal{N} \setminus \{i\}$
7     **else**
8         $v_i = H$
9     Compute the SW-maximising feasible allocation of items to agents in $\mathcal{N}$ and update $I$
10 Return $I$
11 If an agent $i$ is allocated her desired bundle, i.e. $i \in I$, charge her $L$, otherwise charge 0.

$N \setminus I$ has been queried. Mechanism IS belongs to the class $\mathcal{C}$ of mechanisms (Definition 1), and hence by Theorem 2 is IR and OSP. So, we can assume that the agents who have been queried have signaled their true valuation. Regarding the agents that have not been queried, the mechanism assumes valuation $L$ when computing the SW-maximising allocation, while their true valuation could potentially be higher; all these agents have been allocated their desired bundle. This leads to the conclusion that $I$ is the SW-maximising allocation of $\mathcal{N}$ with respect to the true valuations, which in turn implies that $SW(I) \geq SW(I^* \cap \mathcal{N})$.

We now claim that $SW(I) \geq SW(I^* \setminus \mathcal{N})$. Observe that any agent who does not belong to $\mathcal{N}$ has valuation $L$. All these agents were considered in line 2 of Algorithm 1 as having valuation $L$ indeed; denote $SW_1$ the SW computed in line 2 of Algorithm 1. So, it holds that $SW(I^* \setminus \mathcal{N}) \leq SW_1 \leq SW(I)$, where the second inequality holds because the SW computed in Algorithm 1 can only increase between rounds.

Combining the above, we derive that $SW(I^*) = SW(I^* \cap \mathcal{N}) + SW(I^* \setminus \mathcal{N}) \leq 2SW(I)$, as desired. ◀

Next, we show that the above bound is tight by designing a family of instances where no mechanism can achieve an approximation ratio better than 2. This results shows a separation between obviously strategy proof mechanisms and strategy proof mechanisms, in terms of social welfare.

▶ **Theorem 4.** *For the setting that the domain is of size* 2 *and desired bundles are known, no OSP and IR mechanism can achieve an approximation ratio better than* 2*.*

**Proof.** By Theorem 2 we may restrict our analysis to mechanisms in class $\mathcal{C}$ as defined in Definition 1. Consider any such mechanism $M \in \mathcal{C}$.

Consider an instance with $m$ items and a set $N$ of $n = 2\sqrt{m}$ agents. For $i \in \{1, \dots, \sqrt{m}\}$, define sets $S_i = \{(i-1)\sqrt{m}, (i-1)\sqrt{m}+1, \dots, (i-1)\sqrt{m} + (\sqrt{m}-1)\}$ and $T_i = \{i-1, (i-1)+\sqrt{m}, (i-1)+2\sqrt{m}, \dots, i+(\sqrt{m}-1)\sqrt{m}\}$. Let agent $i \in \{1, \dots, \sqrt{m}\}$ desire bundle $S_i$, and let agent $i \in \{\sqrt{m}+1, \dots, 2\sqrt{m}\}$ desire bundle $T_i$. Note that the agents are intuitively split in two equal-sized groups, $A$ and $B$, such that the desired bundle of an agent in $A$

overlaps with the desired bundle of all agents in $B$ and vice versa. Consider such a split and let $G(i) \in \{A, B\}$ denote the group of agents that $i$ belongs to, and $\overline{G(i)} = N \setminus G(i)$. The valuations of the agents in our instance depend on the sequence of queries made in the implementation tree $\mathcal{T}$ of mechanism $M$. We let $H = n/2$ and $L = 1$, and we define the valuations of the divergent agents in turn, starting with the first divergent agent and considering divergent agents in a sequence on the path of $\mathcal{T}$, where previous agents have signaled their true valuation. Let $i$ refer to the agent whose valuation we define at the current step:

**(i)** If agent $i$ is asked an $L$-query, no agent has been asked an $H$-query so far, and at least 3 agents in $\overline{G(i)}$ have not been queried at this point: Set $v_i = L$.

**(ii)** If $i$ is the first agent that is asked an $H$-query, and there are at least 3 agents in $\overline{G(i)}$ that have not been queried at this point: Set $v_i = H$, set $v_j = L$ for each agent $j \in G(i)$ that has not been queried at this point, and set $v_j = H$ for each $j \in \overline{G(i)}$ that has not been queried at this point.

**(iii)** If no agent has been asked an $H$-query so far, and exactly 2 agents in $G(i) \setminus \{i\}$ have not been queried at this point: Set $v_i = H$, and set $v_j = L$ for every agent $j$ that has not been queried at this point.

Observe that the above three points yield a complete and consistent specification of the valuations of all agents.

Clearly, in any mechanism with bounded approximation ratio, one of cases (ii) or (iii) in the definition of valuations above is realised, since otherwise, $M$ will not allocate any desired bundle to any agent. Assume first that case (ii) is realised, and let $i'$ be the agent who is $H$-queried at that point. $i'$ will signal her true valuation $H$, and will be allocated her desired bundle by the definition of an $H$-query. Hence, no agent in $\overline{G(i')}$ can be allocated her desired bundle and there are at least 3 such agents with valuation $H$. Therefore the optimal social welfare in this case is at least $3H + (n/2 - 3)L$, but the mechanism can only allocate to a subset of $G(i)$, resulting in a social welfare of at most $H + (n/2 - 1)L$. Thus, the approximation ratio of $M$ in this case is at least

$$\frac{3H + (n/2 - 3)L}{H + (n/2 - 1)L} = \frac{2n - 3}{n - 1} \tag{2}$$

on this instance.

Suppose now that case (iii) in the definition of valuations above is realised, and let $i^*$ be the agent who is queried at that point. By definition of the instance, the optimal social welfare now is $H + (n/2 - 1)L$. In case $M$ gives the desired bundles to a subset of agents in $G(i^*)$, then this subset may contain only agent $i^*$ and the remaining two agents of $G(i^*)$ who have not been queried yet with valuation $L$, yielding a social welfare of $H + 2L$. Recall that everyone else in $G(i^*)$ has signaled $L$ to an $L$-query, hence can not be allocated her desired bundle. In case $M$ gives the desired bundles to a subset of agents in $\overline{G(i^*)}$, then the social welfare of $M$ is at most $(n/2)L$. Thus, the approximation ratio of $M$ in this case is at least $\frac{H+(n/2-1)L}{\max\{H+2L,(n/2)L\}}$. By setting $H = n/2$ and $L = 1$ we get

$$\frac{H + (n/2 - 1)L}{\max\{H + 2L, (n/2)L\}} = \frac{n - 1}{n/2 + 2} = \frac{2n - 2}{n + 4}. \tag{3}$$

The limit of both bounds (2) and (3) as $n \to \infty$ is 2, and this proves the claim. ◀

We now provide a stronger inapproximability bound for the case where the mechanism is restricted to be *instance-independent*, i.e., the mechanism's implementation tree is not dependent on the demanded bundles of the agents. That is, there is a mapping from nodes in the implementation tree to queried players and respective query types, and this mapping is *fixed*, i.e., it is *not* a function of the demanded bundles of the agents. Note that e.g. the mechanism of Theorem 3 is *not* instance-independent, because the order in which the players are queried and the types of queries that the mechanism asks, do depend on the demanded bundles of the players.

▶ **Theorem 5.** *No instance-independent mechanism has a bounded approximation ratio.*

**Proof.** Consider any instance-independent OSP and IR mechanism $M$. We define the instance (valuations and desired bundles) that yields the upper bound by considering the sequence of queries made in the implementation tree $\mathcal{T}$ of $M$. We start with the first divergent agent and consider divergent agents in a sequence on the path of $\mathcal{T}$, where previous agents have signaled their true valuation. Let $i$ refer to the $i$th queried agent, whose valuation is $v_i$ and whose desired bundle is $R_i$. Let $\overline{Q(i)}$ denote the set of agents who have not been queried yet at the time that $i$ is queried by the mechanism. As $\mathcal{T}$ is instance-independent, we may construct an instance with a bad approximation ratio for this mechanism by letting the agents' valuation and demanded bundles depend on the sequence of queries that the mechanism asks. We do this as follows:

**(i)** If agent $i$ is asked an $L$-query: Set $v_i = L$ and let $R_i$ comprise a single item that does not belong to the desired bundle of any other agent.

**(ii)** If agent $i$ is asked an $H$-query: Set $v_i = H$ and $v_j = H$ for each agent $j \in \overline{Q(i)}$. Let $R_j$ comprise a single item, distinct for each $j \in \overline{Q(i)}$ and not being desired by any previously considered agent. Also, let $R_i = \cup_{j\in\overline{Q(i)}} R_j$.

In any mechanism with a finite approximation ratio, case (ii) in the definition of the instance above is realised, since otherwise, $M$ will not allocate any desired bundle to any agent. Suppose $\ell$ $L$-queries are asked before the first $H$-query on the path of $\mathcal{T}$ discussed in the definition of the instance. The optimal social welfare is at least $\ell L + (n - \ell - 1)H$, while $M$ can only obtain social welfare equal to $H$ by allocating to the agent who got the $H$-query. Thus, the approximation ratio $\rho$ of $M$ is at least

$$\rho \geq \frac{\ell L + (n-\ell-1)H}{H} = \frac{(n-1)H - \ell(H-L)}{H} \geq \frac{(n-1)L}{H}, \tag{4}$$

since $\ell \leq n-1$. The ratio can be made arbitrarily high, for suitable values of $L$ and $H$. ◀

## 3.2 Large domains

In this section we prove an upper bound of $\sqrt{m}$ for the approximation ratio of OSP mechanisms for known-bundle single minded agents and arbitrary domains. We begin with definitions of classes of mechanisms that are of interest. First, we define extremal mechanisms. Informally, the queries of an extremal mechanism always separate an extreme of the current domain of the queried agent from the rest of her current domain (the same extreme is chosen consistently). Formally,

▶ **Definition 6** (Extremal mechanism). *A mechanism with implementation tree $\mathcal{T}$ is an extremal mechanism if for each internal node $u \in \mathcal{T}$, and divergent agent $i = \mathcal{Q}(u)$ at $u$, agent $i$'s current domain $D_i(u)$ is partitioned by the query at $u$ into a singleton, containing the maximum element (or minimum element, consistently), and the remaining elements of $D_i(u)$.*

**Algorithm 2** An extensive-form implementation of the Greedy algorithm in [17] for known-bundle single-minded agents.

1 Define function $\Phi_i$ as $\Phi_i(x) = x/\sqrt{|R_i|}$.
2 $\mathcal{P} \leftarrow \emptyset$ ($\mathcal{P}$ is the set of bundles that have already been allocated) $\mathcal{N} \leftarrow N$ ($\mathcal{N}$ is the set of agents currently under consideration) $\mathcal{D}_i \leftarrow D_i$ for all $i \in N$ ($\mathcal{D}_i$ is the set of values in $i$'s domain currently under consideration)
3 **while** $\mathcal{N} \neq \emptyset$ **do**
4     Let $j = \arg\max_{k \in \mathcal{N}} \Phi_k(\max \mathcal{D}_k)$
5     **if** *there is $S$ in $\mathcal{P}$ such that $R_j \cap S \neq \emptyset$* **then**
6         $\mathcal{N} = \mathcal{N} \setminus \{j\}$
7     **else**
8         Ask $j$ if her valuation is $\max \mathcal{D}_j$
9         **if** yes **then**
10             $\mathcal{P} \leftarrow \mathcal{P} \cup \{R_j\}$
11             $\mathcal{N} = \mathcal{N} \setminus \{j\}$
12         **else**
13             $\mathcal{D}_j = \mathcal{D}_j \setminus \{\max \mathcal{D}_j\}$
14 Return $\mathcal{P}$

▶ **Definition 7** ($\mathcal{C}_d$ mechanisms)**.** *We define a class $\mathcal{C}_d$ of mechanisms, whose implementation tree $\mathcal{T}$ satisfies:*

- *Consider a divergent agent $i$ who is asked to separate between valuations $v_i$ and $v_i'$ with $v_i < v_i'$ at some internal node $u \in \mathcal{T}$. Consider any outcome $o_{v_i}$ in the subtree rooted at $u$ consistent with signaling valuation $v_i$. Then if $i$ is allocated her desired bundle under $o_{v_i}$, she also gets her desired bundle in all possible outcomes consistent with signaling valuation $v_i'$ at $u$.*

▶ **Theorem 8.** *Every extremal mechanism that belongs to class $\mathcal{C}_d$ is OSP.*

Theorem 8 is a special case of Theorem 12 which we present in Section 4. We note that Theorem 8 is one direction of the characterisation in [10].

▶ **Theorem 9.** *Algorithm 2 is an OSP mechanism that achieves an approximation ratio $\rho \leq \sqrt{m}$, for the case of known-bundle single-minded agents.*

**Proof.** The approximation guarantee of the algorithm is well known in the literature, cf. [17]. By Theorem 8, it remains to prove that Algorithm 2 is extremal and belongs to class $\mathcal{C}_d$. Indeed, observe that queries only take place in line 8 of the algorithm, where the corresponding divergent agent is asked to separate the maximum value in her current domain from all other possible values.

We will now prove that the algorithm belongs to class $\mathcal{C}_d$ (see Definition 7). Indeed, when an agent is queried in line 8 then she is allocated her set when replying yes (as feasibility has been previously guaranteed). Since we only query for the maximum in the current domain of each agent, this implies that the agent would still be allocated her desired bundle had her true valuation (and signals) been higher. ◀

## 4 Deterministic mechanisms for unknown-bundle single-minded agents

In this section we prove an upper bound to the approximability of the optimal social welfare of single-minded combinatorial auctions by OSP mechanisms. It is now assumed that both the desired bundle of the agent and her valuation for it belong to her type, hence are private information. We first define a general class of mechanisms which we prove are OSP. We then provide an implementation of the known Greedy-by-valuation mechanism and prove that it belongs to this class.

Let us define the class of valuation-extremal mechanisms. Informally, the queries of a valuation-extremal mechanism at every node separates extreme valuations in the current domain of the queried agent from the rest of her current domain (where the same extreme is chosen consistently). Formally,

▶ **Definition 10** (Valuation-extremal mechanism). *A mechanism with implementation tree $\mathcal{T}$ is a valuation-extremal mechanism if for each internal node $u \in \mathcal{T}$, and divergent agent $i = \mathcal{Q}(u)$ at $u$, it holds that either agent $i$'s current domain, $D_i(u)$, comprises a single valuation, or $D_i(u)$ is partitioned by the query at $u$ into a set containing all the pairs $(v, S) \in D_i(u)$, where $v$ is the maximum valuation (or minimum valuation, consistently) in $D_i(u)$, and the remaining elements of $D_i(u)$.*

▶ **Definition 11** ($\mathcal{C}_d^u$ mechanisms). *We define a class $\mathcal{C}_d^u$ of mechanisms, whose implementation tree $\mathcal{T}$ satisfies:*

- *Consider a divergent agent $i$ who is asked to separate between $(v_i, S_i) \neq (v'_i, S'_i)$ with $v'_i \geq v_i$ and $S' \subseteq S$, at some internal node $u \in \mathcal{T}$. Consider any outcome $o_{(v_i,S_i)}$ in the subtree rooted at $u$ consistent with signalling type $(v_i, S_i)$. Then if $i$ is allocated her desired bundle under $o_{(v_i,S_i)}$, she also gets her desired bundle in all possible outcomes consistent with signalling type $(v'_i, S'_i)$ at $u$.*

▶ **Theorem 12.** *Every valuation-extremal mechanism that belongs to class $\mathcal{C}_d^u$ is OSP.*

Theorem 12 extends the result in [10] to the more general case of multi-dimensional agents. Due to lack of space, the proof is omitted.

Algorithm 3 presents an implementation of the known Greedy-by-valuation mechanism for the case of unknown-bundle single minded agents.

▶ **Theorem 13.** *Algorithm 3 is an OSP mechanism that achieves an approximation ratio $\rho \leq \delta$ for the case of unknown-bundle single-minded agents, where $\delta$ is the size of the largest desired set.*

**Proof.** The approximation guarantee of the algorithm is folklore. By Theorem 12, it remains to prove that Algorithm 3 is valuation-extremal and belongs to class $\mathcal{C}_d^u$. Indeed, note that when an agent is queried in line 4 then she is separating all types with valuation $v'$, which is the maximum compatible valuation at this point of the execution, with all types consistent to smaller values $v$ (regardless of the desired bundles).

We will now prove that the algorithm belongs to class $\mathcal{C}_d^u$ (see Definition 11). Let $u$ be the node of the implementation tree in which $(v, S)$ is separated from $(v', S')$ with $v' \geq v$ and $S' \subseteq S$; clearly there is nothing to separate if $v = v'$ and $S = S'$. Consider first the case that $v' > v$; $u$ corresponds to a query at line 4. Assume that there exists $\mathbf{b}_{-i} \in D_{-i}(u)$ for which $S$ is won by agent $i$ when playing according to $(v, S)$. This means that $S$ is feasible

**Algorithm 3** An extensive-form implementation of the Greedy-by-valuation algorithm.

1 $\mathcal{P} \leftarrow \emptyset$ ($\mathcal{P}$ is the set of bundles that have already been allocated) $\mathcal{N} \leftarrow N$ ($\mathcal{N}$ is the set of agents currently under consideration) $\mathcal{D}_i \leftarrow D_i$ for all $i \in N$ ($\mathcal{D}_i$ is the set of values in $i$'s domain currently under consideration)
2 **while** $\mathcal{N} \neq \emptyset$ **do**
3     Let $j = \arg\max_{k \in \mathcal{N}} \max \mathcal{D}_k$
4     Ask $j$ if her valuation is $\max \mathcal{D}_j$
5     **if** yes **then**
6         Ask agent $j$ to reveal her desired set $R_j$
7         **if** *there is* $S$ *in* $\mathcal{P}$ *such that* $R_j \cap S \neq \emptyset$ **then**
8             $\mathcal{N} = \mathcal{N} \setminus \{j\}$
9         **else**
10             $\mathcal{P} \leftarrow \mathcal{P} \cup \{R_j\}$
11             $\mathcal{N} = \mathcal{N} \setminus \{j\}$
12     **else**
13         $\mathcal{D}_j = \mathcal{D}_j \setminus \{\max \mathcal{D}_j\}$
14 Return $\mathcal{P}$

at $u$; since $S' \subseteq S$ then $S'$ is also feasible at that point. Therefore, the feasibility check in line 7 is successful and $S'$ is won by agent $i$ playing according to $(v', S')$ irrevocably at that point, that is for any $\mathbf{b}_{-i} \in D_{-i}(u)$, as requested. Consider now the case in which $v' = v$ and $S' \subset S$; $u$ corresponds to a query at line 6. Again, if $S$ is feasible at this point of the execution, so is $S'$. Therefore, $S'$ is allocated if $S$ is, and the proof is complete. ◀

## 5 Randomised mechanisms for known-bundle single-minded agents

In this section we consider randomised mechanisms that are universally OSP. We note that the bounded rationality assumption that motivates OSP does not prevent agents from using and understanding such a mechanism, as they do not need to compute expected utilities to determine obvious dominance. We start by presenting a class of randomised mechanisms that are universally OSP.

▶ **Definition 14** ($\mathcal{M}_R$ mechanisms)**.** *We define a class $\mathcal{M}_R$ of randomised mechanisms, that work as follows:*

- *Fix some probability distribution $F$ on the domain $D$ of valuations of all agents. A mechanism $M \in \mathcal{M}_R$ selects one of the values $v$ in $D$ according to $F$ and asks every agent if her valuation is at least equal to $v$. $M$ selects the maximum number of agents whose requests can be satisfied simultaneously, allocates them their desired bundles and charges each of them $v$.*

▶ **Lemma 15.** *Any mechanism $M \in \mathcal{M}_R$ is universally OSP.*

**Proof.** We claim that any mechanism that belongs to class $\mathcal{M}_R$ is a randomisation among different deterministic OSP mechanisms. Indeed, fix a value $v \in D$ and consider the mechanism that queries all agents if their valuation is at least equal to $v$ and outputs a feasible solution that allocates their desired bundles to the maximum number of agents who reply *yes*. It should be easy to see that if, when queried, an agent signals *yes* when her true valuation is smaller than $v$, then the best she can achieve is 0 utility, which is what she would achieve by signalling truthfully. On the other hand, if, when queried, an agent signals *no* when her true valuation is at least equal to $v$, then the best she can achieve is 0 utility, which is at least what she would achieve by signalling truthfully. ◀

▶ **Theorem 16.** *Consider mechanism $M_F \in \mathcal{M}_R$ (Definition 14), where $F$ is a probability distribution assigning probability $p_j$ to value $V_j \in D$, with $V_j > V_{j-1}$, for $j \in \{1, \dots, d\}$ (set $V_0 = 0$), as follows:*

$$p_1 = \left( d - \sum_{j \le d} \frac{V_{j-1}}{V_j} \right)^{-1} \quad \text{and} \quad p_j = p_1 \left( 1 - \frac{V_{j-1}}{V_j} \right), \text{ for } j \in \{2, \dots, d\}.$$

*Then $M_F$ is a universally OSP randomised mechanism that achieves a $\min \left\{ d, 1 + \ln \left( \frac{V_d}{V_1} \right) \right\}$-approximation of the optimal SW for known-bundle single-minded agents, where $d = |D|$.*

**Proof.** Lemma 15 straightforwardly implies that $M_F$ is universally OSP. Denote by $T_j$ the set of agents who are allocated their desired bundles after $M_F$ selects $V_j$ with probability $p_j$. Also, let $T_j^*$ denote the set of agents $i$ such that $v_i = V_j$ who are allocated their desired bundle in the optimal allocation.

For the optimal social welfare, $SW^*$, it holds that

$$SW^* = \sum_{j \le d} V_j \cdot |T_j^*|$$

We bound the social welfare $SW$ of $M_F$ as follows.

$$\begin{aligned} SW &= \sum_{j \le d} p_j V_j \cdot |T_j| \ge \sum_{j \le d} p_j V_j \left( \sum_{j \le \ell \le d} |T_\ell^*| \right) \\ &= \sum_{j \le d} |T_j^*| \cdot \sum_{\ell \le j} p_\ell V_\ell = \sum_{j \le d} |T_j^*| \sum_{\ell \le j} p_1 \left( 1 - \frac{V_{\ell-1}}{V_\ell} \right) V_\ell \\ &= p_1 \sum_{j \le d} |T_j^*| \sum_{\ell \le j} (V_\ell - V_{\ell-1}) = p_1 \sum_{j \le d} |T_j^*| \cdot V_j \\ &= p_1 \cdot SW^*, \end{aligned}$$

where the inequality holds because by definition of mechanism $M_F$, $T_j$ is the maximum cardinality set of agents with valuations at least $V_j$ that can be satisfied simultaneously, and the third equality uses the definition $V_0 = 0$.

Thus, the approximation ratio $\rho$ of this mechanism satisfies

$$\rho \le \frac{1}{p_1} = d - \sum_{2 \le j \le d} \frac{V_{j-1}}{V_j} \tag{5}$$

So, the right-hand-side of (5) is strictly less than $d$. It remains to prove that $\rho \le 1 + \ln(\frac{V_d}{V_1})$. The expression in (5) is maximised if $\sum_{2 \le j \le d} V_{j-1}/V_j$ is minimised. Letting $X_1, \dots, X_{d-1}$ denote the ratios $V_1/V_2, \dots, V_{d-1}/V_d$, and letting $c$ denote the ratio $V_1/V_d$, this amounts to minimising $\sum_{i \in [d-1]} X_i$ subject to $\prod_{i \in [d-1]} X_i = c$. By the inequality of arithmetic and geometric means it holds that $\sum_{i \in [d-1]} X_i \ge n \cdot c^{1/d-1}$ for any solution to this minimisation problem, and furthermore this holds with equality for the solution where $X_i = c^{1/d-1}$ for all

$i \in [d-1]$. We conclude from this that the expression in (5)) is maximised if all the ratios $V_{j-1}/V_j$ in the summation are equal (for a fixed choice of $V_d/V_1$), and this is achieved by setting $V_j = V_d^{\frac{j-1}{d-1}}$ so that $V_{j-1}/V_j = (V_d/V_1)^{-1/(d-1)}$, for all $2 \leq j \leq d$. Inequality (5) then yields

$$\rho \leq d - (d-1)\left(\frac{V_d}{V_1}\right)^{\frac{-1}{d-1}} \leq 1 + \ln\left(\frac{V_d}{V_1}\right), \tag{6}$$

where the last inequality holds because $1 + \ln(V_d/V_1)$ is the limit of the left hand side, and the left hand side is increasing in $d$. This completes the proof. ◀

We note that a similar bound of $O(\ln(r))$ can be achieved in settings where the valuation space is uncountable and contained in an interval $[a, b]$ with $b/a = r$. One can define a continuous version of $M_F$ by deriving its limit running on $D_\epsilon$ as $\epsilon$ approaches 0. The mechanism would then work as follows:

**(i)** Draw a valuation $x \in [a, b]$ according to the cumulative distribution function $F$, with $F(x) = \frac{\ln(x/a)+1}{\ln(r)+1}$.
**(ii)** Ask each agent whether her valuation exceeds $x$.
**(iii)** Compute the maximum cardinality feasible solution $T_x$ among all yes-responding agents.

A similar analysis to that in the proof of Theorem 16 yields the desired bound.

## References


1. Itai Ashlagi and Yannai A. Gonczarowski. Stable matching mechanisms are not obviously strategy-proof. *Journal of Economic Theory*, 177:405–425, 2018.
2. Lawrence M Ausubel. An efficient ascending-bid auction for multiple objects. *American Economic Review*, 94(5):1452–1475, 2004.
3. Sophie Bade and Yannai A. Gonczarowski. Gibbard-satterthwaite success stories and obvious strategyproofness. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, EC '17, page 565, New York, NY, USA, 2017. Association for Computing Machinery. `doi:10.1145/3033274.3085104`.
4. Gary Charness and Dan Levin. The origin of the winner's curse: A laboratory study. *American Economic Journal: Microeconomics*, 1(1):207–36, February 2009. `doi:10.1257/mic.1.1.207`.
5. Paul Dütting, Vasilis Gkatzelis, and Tim Roughgarden. The performance of deferred-acceptance auctions. *Math. Oper. Res.*, 42(4):897–914, 2017.
6. Ignacio Esponda and Emanuel Vespa. Hypothetical thinking and information extraction in the laboratory. *American Economic Journal: Microeconomics*, 6(4):180–202, November 2014. `doi:10.1257/mic.6.4.180`.
7. Diodato Ferraioli, Adrian Meier, Paolo Penna, and Carmine Ventre. On the approximation guarantee of obviously strategyproof mechanisms. *arXiv preprint*, 2018. `arXiv:1805.04190`.
8. Diodato Ferraioli, Adrian Meier, Paolo Penna, and Carmine Ventre. Automated optimal OSP mechanisms for set systems - the case of small domains. In *Web and Internet Economics - 15th International Conference, WINE 2019, New York, NY, USA, December 10-12, 2019, Proceedings*, pages 171–185, 2019.
9. Diodato Ferraioli, Adrian Meier, Paolo Penna, and Carmine Ventre. Obviously strategyproof mechanisms for machine scheduling. In *27th Annual European Symposium on Algorithms, ESA 2019, September 9-11, 2019, Munich/Garching, Germany*, pages 46:1–46:15, 2019.
10. Diodato Ferraioli, Paolo Penna, and Carmine Ventre. Two-way greedy: Algorithms for imperfect rationality. Submitted for publication, 2020.
11. Diodato Ferraioli and Carmine Ventre. Obvious strategyproofness needs monitoring for good approximations. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.

12 Diodato Ferraioli and Carmine Ventre. Probabilistic verification for obviously strategyproof mechanisms. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 240–246. International Joint Conferences on Artificial Intelligence Organization, July 2018.

13 Diodato Ferraioli and Carmine Ventre. Obvious strategyproofness, bounded rationality and approximation. In Dimitris Fotakis and Evangelos Markakis, editors, *Algorithmic Game Theory*, pages 77–91. Springer, 2019.

14 Magnús M. Halldórsson. Approximations of weighted independent set and hereditary subset problems. In *Computing and Combinatorics*, pages 261–270, Berlin, Heidelberg, 1999. Springer.

15 John H. Kagel, Ronald M. Harstad, and Dan Levin. Information impact and allocation rules in auctions with affiliated private values: A laboratory study. *Econometrica*, 55(6):1275–1304, 1987. URL: http://www.jstor.org/stable/1913557.

16 Maria Kyropoulou and Carmine Ventre. Obviously strategyproof mechanisms without money for scheduling. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1574–1581. International Foundation for Autonomous Agents and Multiagent Systems, 2019.

17 Daniel Lehmann, Liadan Ita Oundefinedallaghan, and Yoav Shoham. Truth revelation in approximately efficient combinatorial auctions. *J. ACM*, 49(5):577–602, September 2002. doi:10.1145/585265.585266.

18 Shengwu Li. Obviously strategy-proof mechanisms. *American Economic Review*, 107(11):3257–87, November 2017. doi:10.1257/aer.20160425.

19 Andrew Mackenzie. A revelation principle for obviously strategy-proof implementation. Research Memorandum 014, Maastricht University, Graduate School of Business and Economics (GSBE), May 2018.

20 Marek Pycia and Peter Troyan. Obvious dominance and random priority. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, EC '19, page 1, New York, NY, USA, 2019. Association for Computing Machinery. doi:10.1145/3328526.3329613.