On The Power of Randomization for Obviously Strategy-Proof Mechanisms*

Shiri Ron[1], **Daniel Schoepflin**[2]

[1]Weizmann Institute of Science
[2]Rutgers University - DIMACS
shiriron@weizmann.ac.il, ds2196@dimacs.rutgers.edu

Abstract

We investigate the problem of designing randomized obviously strategyproof (OSP) mechanisms in several canonical auction settings. Obvious strategyproofness, introduced by Li [American Economic Review 2017], strengthens the well-known concept of dominant-strategy incentive compatibility (DSIC). Loosely speaking, it ensures that even agents who struggle with contingent reasoning can identify that their dominant strategy is optimal.

Thus, one would hope to design OSP mechanisms with good approximation guarantees. Unfortunately, Ron [SODA 2024] has showed that deterministic OSP mechanisms fail to achieve an approximation better than the minimum of the number of items and the number of bidders, even for the simple settings of additive and unit-demand bidders. We circumvent these impossibilities by showing that randomized mechanisms that are obviously strategy-proof in the universal sense obtain a constant factor approximation for these classes. We show that this phenomenon occurs also for the setting of a multi-unit auction with single-minded bidders. Thus, our results provide a more positive outlook on the design of OSP mechanisms and exhibit a stark separation between the power of randomized and deterministic OSP mechanisms.

To complement the picture, we provide lower bounds on the performance of randomized OSP mechanisms in each setting. This further demonstrates that OSP mechanisms are significantly weaker than dominant-strategy mechanisms: it is well known that the deterministic VCG mechanism outputs an optimal allocation in dominant-strategies, whereas we show that even randomized OSP mechanisms cannot obtain more than 87.5% of the optimal welfare.

1 Introduction

Economics is the science of how to allocate scarce resources to several competing parties. In particular, auctions serve as a useful playground to understand who should get what and for what price. We assume a good-willed central planner who aims to allocate the resources in a way that maximizes the social welfare of all parties involved. To achieve that, she has to overcome the following obstacle: the information of bidders is private and they are interested in maximizing their own utility. Therefore, she must carefully design the elicitation mechanism to align the incentives of the agents with her own objective.

For a specific example, consider the simplest setting where an auctioneer wants to allocate a single item among a set $\mathcal{N}$ of bidders. Each bidder i has a value v_i for receiving the good. To maximize social welfare, the auctioneer should give the item to the bidder of highest value, i.e., $i^* = \text{argmax}_{i\in\mathcal{N}} v_i$. For that, the auctioneer must design a mechanism which collects information about the value of the bidders and decides which bidder wins the good and the payment $p_i \geq 0$ of each bidder $i \in \mathcal{N}$.

The classic solution proposed for this setting is the *sealed-bid* second-price auction, wherein bidders report their values directly to the auctioneer and the highest-valued bidder is awarded the good at the second highest price (Vickrey 1961). It is well-known that this auction is *dominant-strategy incentive compatible* (i.e., *strategyproof*), meaning that each bidder maximizes her utility by truthfully reporting her private value regardless of the reports of the other bidders. Theory suggests, therefore, that bidders should never misreport their value in this auction. However, in practice, "real-world" bidders report bids not equal to their true value (Kagel, Harstad, and Levin 1987). Thus, there appears to be a mismatch between the prediction of the theory of strategyproof mechanisms and the observed outcomes.

An alternative to the sealed-bid second-price auction for the single-item setting is the ascending price (Japanese) auction. In this auction, a price clock gradually increases over time and bidders drop out whenever the asking price becomes too high. The ascending price auction implements the exact same outcome as the sealed-bid auction: it awards the item to the highest-valued bidder at the second-highest value. However, bidders appear empirically more likely to follow their optimal truthful strategy when facing an ascending-price auction compared to the sealed-bid format.

To address this discrepancy, Li has defined the notion of *obvious strategyproofness* (OSP), a strengthening of strategyproofness (Li 2017). Loosely speaking, if a mechanism is OSP then even agents unable to perform contingent reasoning can identify that truth-telling is the optimal strategy. Obvious strategy-proofness provides a theoretical explanation for the prevalence of the ascending auction, by claiming that its popularity stems from the fact that it is simpler for

*The majority of the proofs and additional results can be found in the full version.

bidders to understand than the sealed-bid format.

Since its introduction by (Li 2017), obvious strategyproofness has emerged as a "gold standard" for strategic simplicity in mechanism design and the notion has attracted a great deal of attention. For instance, various refinements and relaxations of OSP have been proposed (e.g., (Pycia and Troyan 2023b; Nagel and Saitto 2023; Ferraioli and Ventre 2023a)) and the design of OSP mechanisms has been examined various settings, including, e.g., one-sided matching (Troyan 2019) and two-sided matching markets (Ashlagi and Gonczarowski 2018; Thomas 2021), scheduling (Ferraioli et al. 2019), voting problems (Bade and Gonczarowski 2017; Arribillaga, Massó, and Neme 2020) and allocation problems (Bade and Gonczarowski 2017). Another line of work has aimed to find *characterizations* of OSP mechanisms in various domains. For instance, (Li 2017) showed that for all single-parameter binary allocation settings the class of OSP mechanisms coincides with the class of "personalized clock auctions" (essentially, a natural generalization of the Japanese auction for a single item), and further work has established characterizations for all linear single-parameter domains (Ferraioli, Penna, and Ventre 2021; Ferraioli and Ventre 2023b) and beyond (Pycia and Troyan 2023a; Mackenzie 2020). Using these characterizations, both lower bounds (impossibility results) and upper bounds (mechanisms with proven guarantees) have been proposed for various single-parameter auction settings, including binary allocation with general feasibility constraints (Dütting, Gkatzelis, and Roughgarden 2014; Christodoulou, Gkatzelis, and Schoepflin 2022; Feldman et al. 2022) and procurement settings (Balkanski et al. 2022).

Beyond single-parameter settings, the picture regarding the performance of OSP mechanisms in auctions becomes somewhat pessimistic. (Bade and Gonczarowski 2017) initially showed that, for additive bidders, no obviously strategy-proof mechanism optimizes the welfare, and recent work of (Ron 2024) has, essentially, "closed the book" on deterministic OSP mechanisms for multi-parameter combinatorial auction settings. Even if bidders have additive or unit-demand valuation functions (which are commonly thought to be very "easy" settings), the trivial OSP mechanism which runs an ascending-price auction for the grand bundle achieves the *best possible* approximation guarantee of $\min\{m, n\}$, where m is the number of items and n is the number of bidders. To circumvent this strong impossibility, we turn our attention to *randomized* OSP mechanisms.

Our Results Our main results are upper bounds and lower bounds for randomized OSP mechanisms. We focus our attention on "universally" OSP mechanisms, i.e., mechanisms which are a distribution over deterministic OSP mechanisms. We analyze all the settings considered by (Ron 2024) and show that:

1. For additive bidders in a combinatorial auction, there is a mechanism that obtains a 4 approximation and no mechanism has approximation better than $\frac{8}{7} \approx 1.14$ (Theorem 5 and Theorem 6).
2. For unit-demand bidders in a combinatorial auction, there is a mechanism that obtains an $e \approx 2.72$ approximation and no mechanism has approximation better than than $\frac{8}{7} \approx 1.14$ (Theorem 7 and Theorem 8).
3. For single-minded bidders in a multi-unit auction with unknown demands, there is a mechanism that obtains a 400 approximation and no mechanism has approximation better than 1.2 (Theorem 2 and Theorem 4).

All the impossibilities are for mechanisms that satisfy individual rationality and no-negative-transfers. Likewise, our proposed mechanisms conform to these conditions.

Observe that our upper bounds demonstrate the power of randomization for obviously strategy-proof mechanism design: whilst deterministic OSP mechanisms can only obtain an approximation of $\{m, n\}$ to the optimal welfare (Ron 2024), all these classes admit a randomized OSP mechanism that gives a constant factor approximation to the optimal welfare. In addition, we observe that the randomized $poly(m)$-communication mechanisms that are dominant-strategy incentive compatible and obtain the state of the art approximation guarantees for "richer" classes of valuations in combinatorial auctions are in fact obviously strategy-proof (see Claims 4 and 5).[1]

Our upper bounds are motivated by the following observation regarding the lower bound constructions of (Ron 2024): They are based on the fact that every mechanism that provides a non-trivial approximation to the welfare satisfies that the first bidder that "speaks" in the mechanism does not have an obviously dominant strategy. The underlying cause of this phenomenon is that when querying a bidder for the first time, the mechanism fails because it has no information regarding the valuations of the other bidders. Thus, to overcome this impossibility, our proposed mechanisms are based on the classic secretary approach of sampling a sufficient fraction of the bidders and aggregating their information to determine a price per item. Owing to the use of randomization, this can be done in an obviously dominant manner while maintaining a high fraction of the welfare in expectation.

Our lower bounds for combinatorial auctions with unit-demand and additive bidders further emphasize the restrictiveness of obvious strategy-proofness compared to implementation in dominant strategies. Not only getting an approximation better than $\min\{m, n\}$ is impossible deterministically, but even if we allow randomization we cannot get more than 87.5% of the optimal welfare. In contrast, these settings have dominant-strategy mechanisms that extract the optimal welfare and are also efficient both from a computation and a communication perspective.

One disadvantage of our main results is that the lower bounds and upper bounds that we provide are quite far apart. As a step to bridge this gap, we show in the full version that for two bidders and two items, all the aforementioned classes admit mechanisms that give a $\frac{4}{3}$ approximation.

[1] We also provide a 400 approximation to the optimal welfare for multi-unit auctions with bidders whose valuations satisfy decreasing marginal utilities (Theorem 3). This is the only multi-parameter domain for which the power of deterministic mechanisms is not known. We provide additional results in the full version which explain the difficulty of proving lower bounds for this class.

Why Randomization? On first impression, one might argue that randomization adds impractical complexity to a mechanism, and does not align with the simplicity we aim to achieve when designing OSP mechanisms. Indeed, "real-world" agents could possibly be confused by randomization. Moreover, it can be difficult to verify that an outcome is the result of some pre-specified random process. We emphasize, however, that the randomization at use in our work is rather "straightforward", and, crucially, as we discussed above, bidders in our auctions need not reason about expected outcomes since, on any fixed result of the random process, they face a mechanism where they have an obviously dominant strategy. Moreover, randomized mechanisms are prevalent in practice, e.g., in drafts in sports for new team members, housing programs, and "greencard" allocation, and well-studied elsewhere in theory, e.g., in fair division (see (Budish et al. 2013) among many others). In fact, it is not clear how to design fair mechanisms without randomization, as the order of selection of players often plays a crucial role. Randomizing the order of choice is then a natural solution, and it is unclear whether there are any alternatives to this approach.

The power of randomization for welfare maximization is, despite a great deal of work, still not fully understood. If we put aside complexity considerations, then the deterministic VCG mechanism is optimal, so randomized dominant-strategy mechanisms do not "beat" their randomized counterparts. In contrast, if we require communication-efficient mechanisms, then randomized dominant-strategy mechanisms achieve an approximation of $\mathcal{O}(\sqrt{m})$ when the bidders have arbitrary monotone valuations, whilst every deterministic mechanism that achieves approximation better than $\mathcal{O}(m^{1-\epsilon})$ has exponential communication (Dobzinski, Nisan, and Schapira 2012; Dobzinski, Ron, and Vondrák 2022). If we settle for mechanisms that only satisfy the weaker notion of ex-post incentive compatibility[2], then the best deterministic mechanisms lag behind their randomized counterparts (Qiu and Weinberg 2024; Assadi and Singla 2019; Assadi, Kesselheim, and Singla 2021), but no such separation is known. Our work, in contrast, demonstrates a separation between deterministic and randomized mechanisms for welfare maximization if we strengthen the incentive property to OSP.

Discussion and Open Questions For deterministic OSP mechanisms, ascending auctions, namely (deferred acceptance) clock auctions, seem to be the "holy grail". Unfortunately, they seem to do poorly. Using randomization, however, we bypass the impossibilities for deterministic OSP mechanisms. While our "lottery over ascending auctions" only guarantees a logarithmic-approximation for, e.g., multi-unit auctions, our random sampling-based mechanisms[3] give us $\mathcal{O}(1)$-approximations in each setting we explore.

We note that there is a gap between our upper and lower bounds and we leave open the question of understanding the exact approximation ratio of these classes of mechanisms and of randomized OSP mechanisms in general. A particular enticing question is whether there exists an $\mathcal{O}(1)$-approximate randomized OSP mechanism for more general complement-free valuations (e.g., submodular or subadditive valuations) or whether one can demonstrate an $\omega(1)$-lower bound on the performance of any OSP mechanism in these settings.

Toward the latter, one would need to circumvent a crucial limitation of our lower bound approach. The reason for it is that we prove our lower bounds by presenting a distribution over valuation profiles, and showing that a mechanism has to fail on at least one of them. However, in order to prove super-constant lower bounds, it is essential to find distributions whose support has super-constant number of instances and show that every deterministic mechanism errs on a significant fraction of them.

There are also interesting questions outside the realm of the auction settings. For example, exploring other settings to determine whether there exists a separation between the performance dominant-strategy and obviously strategy-proof mechanisms is an intriguing direction.

2 Preliminaries

We consider settings where an auctioneer aims to allocate a set M of m items to a set $\mathcal{N}$ of n bidders. To do so, the auctioneer designs randomized mechanisms. A randomized protocol $\mathcal{M}$ randomly chooses, in advance, one of several deterministic protocols to follow. We denote the deterministic mechanisms in the support of the randomized mechanism with $\mathcal{A}$. Throughout the paper, we use the terms protocol and mechanism interchangeably.

Components of Deterministic Protocols Following the random selection, all bidders face a deterministic protocol. We represent protocols as trees where each internal node corresponds to a bidder called to "speak" or communicate a message. Each such node has a set of possible messages, and the next node in the tree is determined by the message sent by the bidder. A leaf in the protocol specifies an allocation of items to bidders and a payment for each bidder.

Fixing a randomly chosen deterministic protocol $A \in \mathcal{A}$, let $\mathcal{N}_i$ denote the set of all nodes in which a particular bidder i is called to speak. Then, the behavior B_i of player i assigns a message to each node in $\mathcal{N}_i$. We denote with $\mathcal{B}_i$ be the set of all possible behaviors. Observe that a behavior profile $B = (B_1, \dots, B_n) \in \mathcal{B}_1 \times \dots \times \mathcal{B}_n$ thus defines a root to leaf path in A. We let $Path(B)$ denote all the nodes along the path defined by B and $Leaf(B)$ denote the leaf that $Path(B)$ ends with. For every behavior profile B and for every player i, we denote with $f_i(B)$ and with $p_i(B)$ respectively the allocation and the payment of player i that are specified in $Leaf(B)$. Finally, the *strategy* $\mathcal{S}_i$ of player i is a function specifying a behavior of player i for each possible valuation in V_i and every possible protocol A. Formally, $\mathcal{S}_i : \mathcal{A} \times V_i \to \mathcal{B}_i$. We often abuse notation by also naming the strategy of player i in the deterministic protocol A, i.e.,

[2] A mechanism is ex-post incentive compatible if it has strategies that form a Nash equilibrium, in contrast to dominant-strategy mechanisms where the bidders have dominant strategies.

[3] These auctions are not expressible as randomized deferred acceptance clock auctions because we allow bidders to select their favorite bundles of goods at a given price vector.

the partial function $\mathcal{S}_i(A, \cdot)$, also a strategy and denoting it with $\mathcal{S}_i$.

We say that a deterministic mechanism A together with strategies $(\mathcal{S}_1(A, \cdot), \ldots, \mathcal{S}_n(A, \cdot))$ realize allocation rule $f : V_1 \times \cdots \times V_n \to \mathcal{T}$ and payment schemes $P_1, \ldots, P_n : V_1 \times \cdots \times V_n \to \mathbb{R}^n$ if for every $(v_1, \ldots, v_n) \in V_1 \times \cdots \times V_n$, it holds that $Leaf(B_1, \ldots, B_n)$ is labeled with the allocation $f(v_1, \ldots, v_n)$ and with the payment $P_i(v_1, \ldots, v_n)$ for every player i, where $\forall i \in \mathcal{N}$ the behavior $B_i = \mathcal{S}_i(A, v_i)$.

Properties of Randomized Mechanisms To analyze the performance of our mechanisms, we compare against the optimal social welfare. Let $\mathbf{T} = (T_1, \ldots, T_n)$ denote a feasible allocation of the items (i.e., each item is allocated to at most one bidder) and $\mathcal{T}$ denote the set of all feasible allocations. Then, we let $\text{OPT}(I) = \max_{\mathbf{T} \in \mathcal{T}} \sum_{i \in \mathcal{N}} v_i(T_i)$ denote the optimal social welfare achievable on a given instance $I = (v_1, \ldots, v_n)$ and $\mathbb{E}[W(\mathcal{M}(I))]$ denote the *expected* social welfare achieved by mechanism $\mathcal{M}$ on instance I (where the expectation is taken over the random choice of which deterministic protocol is to be run by the mechanism). We then say that mechanism $\mathcal{M}$ obtains an α-approximation to the optimal social welfare on a class of instances $\mathcal{I}$ if

$$\sup_{I \in \mathcal{I}} \frac{\text{OPT}(I)}{\mathbb{E}[W(\mathcal{M}(I))]} \leq \alpha.$$

In addition to the objective of welfare maximization, our goal is to design randomized mechanisms satisfying three key desiderata: (i) *no-negative transfers*; (ii) *ex-post individual rationality*; and (iii) *universal obvious strategyproofness*, i.e., ex-post obvious strategyproofness.[4] A mechanism $\mathcal{M}$ with support $\mathcal{A}$ satisfies *no-negative transfers* if for every leaf in every deterministic protocol $A \in \mathcal{A}$, the payment of every player i is at least zero. A mechanism $\mathcal{M}$ with strategy profile $(\mathcal{S}_1, \ldots, \mathcal{S}_n)$ and support $\mathcal{A}$ satisfies *ex-post individual rationality* if, for every deterministic protocol $A \in \mathcal{A}$ that realizes an allocation rule f and payment schemes $P_1, \ldots, P_n$, it holds that for every $(v_1, \ldots, v_n)$ and every player i: $v_i(f(v_1, \ldots, v_n)) - P_i(v_1, \ldots, v_n) \geq 0$. Namely, a mechanism is individually rational if each player obtains non-negative utility for participating in the mechanism (hence, there is no incentive to avoid participation).

Finally, we define obvious strategy-proofness in the universal sense. A mechanism $\mathcal{M}$ with support $\mathcal{A}$ and the strategies $(\mathcal{S}_1, \ldots, \mathcal{S}_n)$ is universally obviously strategy-proof if for every deterministic protocol $A \in \mathcal{A}$, the strategies $(\mathcal{S}_1(A, \cdot), \ldots, \mathcal{S}_n(A, \cdot))$ are obviously dominant. It remains to define what it means for a strategy to be obviously dominant. Loosely speaking, a strategy $\mathcal{S}_i(A, \cdot)$ of bidder i in a deterministic mechanism A is *obviously dominant* if, each time player i is called to speak, the worst-case outcome from sending the message defined by $\mathcal{S}_i(A, \cdot)$ is weakly better than the best-case outcome from any other strategy. Thus, it is conceptually "easy" for player i to find $\mathcal{S}_i$, understand its dominance and follow it. Despite their intuitive appeal, the definition of obviously dominant strategies is quite subtle. Thus, we defer the precise definition to the full version.

[4]A randomized mechanism satisfies a given property *ex-post* if that property holds for every deterministic mechanism that has a non-zero probability of being selected.

Generalized Ascending Auctions To prove some of our positive results, we employ a specific form of auction, which we name generalized ascending auctions. In particular, some of our randomized mechanisms will be a randomization of such auctions. A *generalized ascending auction* is an auction that defines two possible allocations for each bidder i, the *base* bundle X_i^B and the *potential* bundle X_i^P where $X_i^B \subseteq X_i^P$. Each bidder i initially "holds" bundle X_i^B and is placed in the "active" set $\mathcal{N}^A$. Each $i \in \mathcal{N}^A$ faces a monotonically increasing price trajectory to instead receive X_i^P (in place of X_i^B). Bidders then drop out when the price for receiving X_i^P is too high at which point they are awarded X_i^B at a price of 0 and removed from $\mathcal{N}^A$ (i.e., become inactive). Finally, the auction terminates when allocating all $i \in \mathcal{N}^A$ their potential bundle X_i^P and all $i \notin \mathcal{N}^A$ their base bundle X_i^B is feasible. For illustration, an auction where bidder 1 always gets some item a and there is an ascending auction on the remaining items $M \setminus \{a\}$ is a generalized ascending auction.

Lemma 1. *Every generalized ascending auction is obviously strategy-proof.*

We defer the proof of Lemma 1 to the full version.

3 Multi-Unit Auctions

We begin with the setting of multi-unit auctions, i.e., an auction with m *identical* items. We first consider the case of *unknown single-minded bidders*, where each bidder i receives value v_i for receiving at least d_i items and 0 otherwise.[5]

Upper Bound: a First Attempt

Since an ascending-price auction for the grand bundle of goods is the optimal deterministic OSP mechanism in this setting (Ron 2024), we begin with a natural randomized analogue of this approach. Namely, we "guess" a bundle size and run an ascending price auction for bundles of this size.

Formally, let $k = \lceil \log m \rceil$. Consider the mechanism RANDOM-BUNDLES in which an integer bundle size $\ell = 2^j$ is sampled uniformly at random from the set $\ell \in \{1, 2, 4, 8, \ldots, \frac{m}{2}, m\}$. Given this fixed bundle size, every bidder either wins exactly ℓ items or wins no items at all. Observe, then, that at most m/ℓ bidders can win a bundle. We now increase the price of being served ℓ items, until at most m/ℓ bidders remain. All these bidders win ℓ items and pay the price at which we stop, while the remainder get nothing and pay nothing. Note that RANDOM-BUNDLES is a generalized ascending auction, so by Lemma 1 it is OSP.

We now argue that this mechanism obtains a poly-logarithmic approximation to the optimal social welfare. Due to space constraints, we provide a proof sketch of the approximation guarantee and defer the complete proof to the full version.

[5]This setting is commonly known as unknown single-minded bidders, since both the demand d_i and the value v_i of every player i are private. If the demand d_i is not private information, then it is a setting with *known single-minded bidders*.

Theorem 1. RANDOM-BUNDLES *obtains an* $\mathcal{O}(\log m)$ *approximation to the optimal social welfare.*

Proof Sketch. First, observe that we can partition bidders into groups depending on their demand as follows: for each bidder i, we place bidder i in group p if her demand d_i is between 2^p and $2^{p+1}-1$ (inclusive). Since each bidder has demand at least 1 and at most m, there are at most $\log m$ groups in total.

Now, we compare the portion of the optimal social welfare coming from bidders in group p against the welfare RANDOM-BUNDLES obtains when selecting bundles of size 2^{p+1}. On one hand, the optimal solution selects at most twice as many bidders appearing in group p as the total number of bidders RANDOM-BUNDLES serves conditioned on it selecting bundles of size 2^{p+1}. On the other hand, the bidders served in RANDOM-BUNDLES conditioned on selecting bundles of size 2^{p+1} have the *highest* value among bidders satisfied by receiving 2^{p+1} goods. In total, we obtain an $\mathcal{O}(\log m)$-approximation. $\square$

A Constant Upper Bound for Multi-Unit Auctions

Unfortunately, an approximation of $\mathcal{O}(\log m)$ seems to be roughly the limit of the approach of randomly choosing fixed bundle sizes. As such, we need to turn to a new approach. We, thus, adapt the "balanced sampling" approach utilized extensively in other areas of mechanism design (see, e.g., (Feige et al. 2005; Goldberg et al. 2006; Dobzinski, Nisan, and Schapira 2012; Dobzinski 2007; Badanidiyuru, Kleinberg, and Singer 2012; Bei et al. 2017)) in the form of Mechanism 1, below:

Theorem 2. *There is a universally OSP mechanism for unknown single-minded bidders in a multi-unit auction that gives a* 400*-approximation to the optimal welfare.*

We prove Theorem 2 by describing a randomized mechanism, i.e., Mechanism 1, and showing that is universally OSP (Lemma 2) and indeed gives a 400-approximation to the optimal social welfare (Lemma 4).

MECHANISM 1: "SINGLE-MINDED"

Input: A set of bidders N and m identical items
With probability $1/2$:
  Bundle all m items together and run an ascending price auction on the grand bundle
With remaining probability $1/2$:
  let $S \leftarrow \emptyset$, $U \leftarrow \emptyset$
  Place each bidder independently in S w.p. $1/2$ and each bidder in U with the remaining probability
  "Discard" each bidder and S and learn their valuation function
  Compute the optimal solution among bidders only in S and let O denote the value of this solution
    Iterate over the bidders (in an arbitrary) order, and for each bidder $i \in \mathcal{N}$, let them purchase their preferred bundle from the remaining items at a price of $O/10m$ per item

Lemma 2. *Mechanism 1 is universally OSP.*

Proof. Under the realization of randomness where we auction the grand bundle, we utilize a generalized ascending auction for the grand bundle which is OSP by Lemma 1.

Consider a fixed realization of randomness where we run the uniform pricing auction. No bidder in S can win any items and, thus, the mechanism is OSP for them. Bidders in U select their favorite bundle of goods and, thus, the auction is OSP for them as well. $\square$

To prove the approximation factor of Mechanism 1, we define a bidder as *critical* if her value for the grand bundle of goods is at least $1/100$ of the total optimal social welfare. We also use the notation $OPT(S)$, where S is a subset of bidders $S \subseteq \mathcal{N}$, to denote the optimal welfare achieved by allocating all items exclusively among the bidders in S.

Lemma 3 establishes that the sampling phase yields "representative" sampled and unsampled sets in the case that there are no critical bidders with "high enough" probability. Note that the proof utilizes a lemma of (Bei et al. 2017). We defer the proof of Lemma 3 to the full version.

Lemma 3. *Consider an instance* $(v_1, \ldots, v_n)$ *of bidders with single-minded valuations*[6] *where no bidder is critical. Suppose each bidder is placed in a "sampled" set* S *with probability* $1/2$ *and placed in an "unsampled" set* U *with the remaining probability independently. Then the optimal welfare obtained from bidders in the sampled set OPT*(S) *and the optimal welfare obtained from bidders in the unsampled set OPT*(U) *are such that OPT*$(S) \geq OPT/5$ *and OPT*$(U) \geq OPT/5$ *with probability at least* $1/2$*.*

With Lemma 3 in hand, we are ready to prove that:

Lemma 4. *Mechanism 1 obtains a* 400*-approximation to the optimal social welfare.*

Proof. First we handle the case that there is a critical bidder. The existence of a critical bidder i implies that allocating i the grand bundle gives a $1/100$-approximation to the optimal welfare. Since we run an ascending auction on the grand bundle with probability $1/2$ we obtain a $1/200$-approximation when there exists a critical bidder.

We now turn to the case that there does not exist a critical bidder. In this case, by Lemma 3, with probability $1/2$ over the random sampling of bidders, the optimal welfare achievable by the sampled set is within a factor 5 of the optimal welfare. As such when we proceed to the pricing phase, we set a price per item $p \in [\text{OPT}/50m, \text{OPT}/10m]$. Since we run a uniform price auction with probability $1/2$, we have that these conditions hold with probability at least $1/4$. We perform case analysis on the number of items sold during this phase.

Suppose our uniform pricing phase sells at least $m/2$ goods. In this case, since an unsampled bidder buying t goods spends at least $\frac{t\text{OPT}}{50m}$, their value for the purchased bundle is at least $\frac{t\text{OPT}}{50m}$. Then, the total value of all bidders who purchase goods is at least $\frac{m}{2} \cdot \frac{\text{OPT}}{50m} = \frac{\text{OPT}}{100}$. Altogether, since we run uniform sampling with probability $1/2$ and the

[6]Our lemma actually holds for more general valuations, but for simplicity we state it for single-minded bidders.

estimation is "good" with probability $1/2$, we obtain a 400-approximation to the welfare.

Now, we analyze the complementary case where the uniform pricing phase sells fewer than $\frac{m}{2}$ goods. For that, let $\vec{q} = (q_1, \ldots, q_n)$ be the optimal allocation if the items are divided only among the bidders in U (clearly, every bidder not in U is allocated zero items).

For that, observe that if a bidder is allocated in $\vec{q}$ but not allocated in the allocation of the algorithm, it happens because of one of the following reasons. The first possibility is that the bidder is *blocked*, meaning that the number of items that she wants d_i is not available when it is her turn. The other reason is that the bidder is *small*, meaning that $v_i(q_i) \le p \cdot q_i$, i.e., the price set is too high for her.

Note that every bidder i that is satisfied (i.e., allocated enough goods to match her demand) in the allocation $\vec{q}$ and is neither blocked nor small, is also satisfied in the algorithm. We will bound the loss of welfare from both kinds of bidders, assuming that we run a uniform price auction and that our sampling was "balanced" (i.e., both $OPT(S) \ge OPT/5$ and $OPT(U) \ge OPT/5$) which, by Lemma 3 occurs with probability at least $1/4$.

First, we show that blocked bidders do not exist, so they do not cause any loss of welfare. First, we remind that by assumption the uniform phase sells less than $\frac{m}{2}$ items. Thus, a blocked bidder wants to purchase strictly more than $\frac{m}{2}$ goods at a price of at least $\frac{\text{OPT}}{50m}$ and thus has a value of at least $\frac{\text{OPT}}{100}$. Since by assumption there are no critical bidders, we get that there are no blocked bidders.

Now, since the sum of items allocated to small bidders in total is at most m, we have that their total welfare in the optimal allocation is at most $\frac{OPT}{10}$. Since the welfare of $\vec{q}$ is at least $\frac{OPT}{5}$, it implies that bidders who are neither blocked or small contribute at least $\frac{OPT}{10}$ to the welfare of $\vec{q}$. Since Mechanism 1 allocates to these bidders their desired number of items, it achieves welfare of at least $\frac{OPT}{10}$. As we said before, this happens with probability $\frac{1}{4}$, so overall the expected welfare of the mechanism is at least $\frac{OPT}{40}$.

Combining all cases, we conclude that the expected welfare of Mechanism 1 is at least $\frac{OPT}{400}$, thereby completing the proof. $\square$

Mechanism 1 also achieves a constant-approximation in the case that bidders have *decreasing marginal valuations*. This is an improvement over the best-known deterministic OSP mechanism of (Gkatzelis, Markakis, and Roughgarden 2017) which achieved only an $O(\log n)$ approximation.

Theorem 3. *Mechanism 1 obtains a* 400*-approximation to the optimal social welfare in the case of bidders with decreasing marginal valuations.*

We defer the proof of Theorem 3 to the full version.

Lower Bound

Theorem 4. *For a multi-unit auction with* $m \ge 2$ *items and* $n \ge 2$ *unknown single-minded bidders, no randomized mechanism that satisfies ex-post OSP, individual rationality and no-negative-transfers has approximation better than* $6/5$.

Proof. We assume the domain V_i of each bidder consists of single-minded valuations with values in $\{0, 1, \ldots, k^4\}$, where k is an arbitrarily large number. Our example has only two bidders, but it can be extended to any number of bidders by adding bidders with the all-zero valuation.

To use our variant of Yao's principle (which is stated in the full version), we define a distribution $\mathcal{D}$ of valuation profiles and show that no deterministic mechanism that satisfies obvious strategy-proofness, individual rationality and no-negative-transfers with respect to $V = V_1 \times V_2$ has approximation better than $\frac{6}{5}$ in expectation over $\mathcal{D}$. To define it, consider the following valuations:

$$v_i^{\text{one}}(x) = \begin{cases} 1 & x \ge 1, \\ 0 & \text{else} \end{cases} \quad v_i^{\text{ONE}}(x) = \begin{cases} k^2+1 & x \ge 1, \\ 0 & \text{else} \end{cases}$$

$$v_i^{\text{all}}(x) = \begin{cases} k^2 & x = m, \\ 0 & \text{else} \end{cases} \quad v_i^{\text{ALL}}(x) = \begin{cases} k^4 & x = m, \\ 0 & \text{else} \end{cases}$$

Consider the following valuation profiles:

$$I_1 = (v_1^{\text{one}}, v_2^{\text{one}}) \quad I_2 = (v_1^{\text{all}}, v_2^{\text{one}}) \quad I_3 = (v_1^{\text{ONE}}, v_2^{\text{ALL}})$$
$$I_4 = (v_1^{\text{one}}, v_2^{\text{all}}) \quad I_5 = (v_1^{\text{ALL}}, v_2^{\text{ONE}})$$

Denote with $\mathcal{D}$ the distribution over profiles where the probability of I_1 is $\frac{1}{3}$, and the probability of I_2, I_3, I_4 and I_5 is $\frac{1}{6}$ each. Observe that:

Claim 1. *Every deterministic mechanism that has approximation strictly better than* $6/5$ *necessarily satisfies all of the following conditions:*

1. *Given the valuation profile* $I_1 = (v_1^{one}, v_2^{one})$*, the mechanism allocates at least one item to every bidder.*
2. *Given the valuation profile* $I_2 = (v_1^{all}, v_2^{one})$*, the mechanism allocates all items to bidder* 1.
3. *Given the valuation profile* $I_3 = (v_1^{ONE}, v_2^{ALL})$*, the mechanism allocates all items to bidder* 2.
4. *Given the valuation profile* $I_4 = (v_1^{one}, v_2^{all})$*, the mechanism allocates all items to bidder* 2.
5. *Given the valuation profile* $I_5 = (v_1^{ALL}, v_2^{ONE})$*, the mechanism allocates all items to bidder* 1.

The proof of Claim 1 is straightforward: if a deterministic mechanism A violates one of the conditions, then since k is arbitrarily large, the mechanism A extracts at most $\frac{5}{6}$ of the optimal welfare in expectation over the distribution $\mathcal{D}$.

Fix a deterministic mechanism A and strategies $(\mathcal{S}_1, \mathcal{S}_2)$ that are individually rational and satisfy no-negative-transfers with respect to the valuations $V_1 \times V_2$ and give approximation better than $\frac{6}{5}$ in expectation over the valuation profiles in the distribution $\mathcal{D}$. Let (f, P_1, P_2) be the allocation and payment rules that the mechanism A and the strategies $(\mathcal{S}_1, \mathcal{S}_2)$ jointly realize. Assume towards a contradiction that A and $(\mathcal{S}_1, \mathcal{S}_2)$ are OSP.

To analyze the mechanism, we focus on the following subsets of the domains of the valuations: $\mathcal{V}_1 = \{v_1^{one}, v_1^{ONE}, v_1^{ALL}\}$ and $\mathcal{V}_2 = \{v_2^{one}, v_2^{ONE}, v_2^{ALL}\}$.[7] Observe that either bidder 1 or 2 has to send different messages

[7]The cautious reader may have noticed that $\mathcal{V}_i$ does not contain v_i^{all}. This is intentional, and it will be clear from the remainder of the proof why including this valuation is not necessary.

for different valuations in $\mathcal{V}_i$ at some vertex. This is an immediate implication of Claim 1, as the mechanism A necessarily outputs different allocations given the valuation profiles $I_1 = (v_1^{one}, v_2^{one})$ and $I_2 = (v_1^{all}, v_2^{one})$, meaning that the behaviors $(\mathcal{S}_1(v_1^{one}, \mathcal{S}_2(v_2^{one}))$ and $(\mathcal{S}_1(v_1^{all}, \mathcal{S}_2(v_2^{one}))$ reach different leaves and thus have to diverge at some point.

Let u be the first vertex in the protocol such that $(\mathcal{S}_1(v_1), \mathcal{S}_2(v_2))$ and $(\mathcal{S}_1(v_1'), \mathcal{S}_2(v_2'))$ diverge, i.e., dictate different messages. Note that by definition $u \in Path(\mathcal{S}_1(v_1), \mathcal{S}_2(v_2)) \cap Path(\mathcal{S}_1(v_1'), \mathcal{S}_2(v_2'))$. We remind that each vertex is associated with only one player that sends messages in it. Note that the distribution $\mathcal{D}$ we have defined is symmetric, so we can assume without loss of generality that player 1 is the player that sends a message in vertex u. Thus, there exist $v_1, v_1' \in \mathcal{V}_1$ such that $\mathcal{S}_1(v_1)$ and $\mathcal{S}_1(v_1')$ dictate different messages at vertex u. We remind that $\mathcal{V}_1 = \{v_1^{one}, v_1^{ONE}, v_1^{all}\}$, so the following claims jointly imply a contradiction, completing the proof:

Claim 2. *The strategy $\mathcal{S}_1$ dictates the same message at vertex u for the valuations v_1^{one} and v_1^{ONE}.*

Claim 3. *The strategy $\mathcal{S}_1$ dictates the same message at vertex u for the valuations v_1^{ONE} and v_1^{ALL}.*

We defer the proofs of Claims 2 and 3 to the full version. The proofs are are based on the properties of the mechanism: its approximation guarantee, obvious strategy-proofness, individual rationality and no-negative-transfers. □

4 Combinatorial Auctions

We now turn toward settings with heterogeneous items. We explore settings involving additive and unit-demand bidders, and wrap up by considering mechanisms for subadditive and general valuations. The proofs of all theorems can be found in the full version.

Additive Valuations

A valuation v_i is *additive* if bidder i has a value $v_{ij} \geq 0$ for item j and the value bidder i has for receiving a set of items A_i is equal to $\sum_{j \in A_i} v_{ij}$.

Upper Bound We show that the sampling approach yields a 4-approximation for this setting:

Theorem 5. *There is a universally OSP mechanism for bidders with additive valuations that gives a 4-approximation to the optimal welfare.*

We provide the description of 2 and defer the correctness proof to the full version due to lack of space. Simply put, Mechanism 2 samples a threshold price for each item, and then uses these threshold prices as a posted-price mechanism for the unsampled bidders. We give priority to high index bidders to handle tie-breaking issues.

Lower Bound

Theorem 6. *Even for two bidders and two items, there is no randomized obviously strategy-proof mechanism that satisfies individual rationality and no-negative-transfers and gives approximation better than $\frac{8}{7}$ to the optimal social welfare for additive bidders.*

MECHANISM 2: "ADDITIVE"

Input: A set of bidders $\mathcal{N}$ and a set of M items
1. Index the bidders in some arbitrary fixed order
2. $S \leftarrow \emptyset, U \leftarrow \emptyset$
3. Independently assign each bidder to set S with probability $1/2$ and to set U with probability $1/2$.
4. "Discard" each bidder in S and learn their value for each item
5. For each $j \in M$: set a price p_j on item j equal to $\max_{i \in S} v_{ij}$ and let $n(j)$ denote the smallest index among bidders in $\arg\max_{i \in S} v_{ij}$
6. For each $i \in U$ in an arbitrary order: Let i purchase all previously unsold items $j \in M$ for which either: (i) $v_{ij} > p_j$; or (ii) $v_{ij} = p_j$ and i has a lower index than $n(j)$

Unit-Demand Valuations

We now address bidders with *unit-demand* valuations, where there is a value $v_{ij} \geq 0$ for each $i \in \mathcal{N}$ and $j \in M$ and the value bidder i has for a set A_i is equal to $\max_{j \in A_i} v_{ij}$.

Upper Bound This setting appears more complicated than the setting of additive valuations, as the approach of setting the price of each item to be the price of the second highest bid fails miserably (we provide an example that demonstrates it in the full version.) Thus, to obtain a mechanism with a constant approximation for this setting, we use the beautiful algorithm of (Reiffenhauser 2019), originally formulated for the problem of strategy-proof online matching:

Theorem 7. *There exists a universally OSP mechanism which achieves an e-approximation to the optimal social welfare in the presence of unit-demand bidders.*

Lower Bound

Theorem 8. *Even for two bidders and two items, there is no randomized obviously strategy-proof mechanism that satisfies individual rationality and no-negative-transfers and gives approximation better than $\frac{8}{7}$ to the optimal social welfare for unit-demand bidders.*

More General Valuations

In light of our previous results, one may wonder if there is any "rich enough" class of valuations for which it is known that randomized OSP mechanisms cannot extract more than a constant fraction of the welfare. Perhaps surprisingly, the answer to this is no. However, we were able to verify that the state-of-the-art *computationally efficient* randomized mechanisms for subadditive and general valuations[8] are, in fact, universally OSP:

Claim 4. *The $O((\log\log(m))^3)$-approximate randomized mechanism of (Assadi, Kesselheim, and Singla 2021) for subadditive valuations is universally OSP.*

Claim 5. *The $O(\sqrt{m})$-approximate randomized mechanism of (Dobzinski, Nisan, and Schapira 2012) for general valuations is universally OSP.*

[8]A function v is subadditive if for every two bundles of items $A, B \subseteq M$, it holds that $v(A \cup B) \leq v(A) + v(B)$.

Acknowledgments

The authors would like to thank Shahar Dobzinski for valuable discussions during the early stages of this project and to Simon Mauras for introducing us to (Reiffenhauser 2019), which was very helpful to our work.

The first author is supported by an Azrieli Foundation fellowship, ISF grant 2185/19, and BSF-NSF grant (BSF number: 2021655, NSF number: 2127781). The second author is supported in part by a grant to DIMACS from the Simons Foundation (820931). Part of this work was done while both authors were in residence at the Simons Laufer Mathematical Sciences Institute, supported by the NSF under Grant No. DMS-1928930 and by the Alfred P. Sloan Foundation under grant G-2021-16778.

References

Arribillaga, R. P.; Massó, J.; and Neme, A. 2020. On obvious strategy-proofness and single-peakedness. *Journal of Economic Theory*, 186: 104992.

Ashlagi, I.; and Gonczarowski, Y. A. 2018. Stable matching mechanisms are not obviously strategy-proof. *Journal of Economic Theory*, 177: 405–425.

Assadi, S.; Kesselheim, T.; and Singla, S. 2021. Improved truthful mechanisms for subadditive combinatorial auctions: Breaking the logarithmic barrier. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 653–661. SIAM.

Assadi, S.; and Singla, S. 2019. Exponentially Improved Truthful Combinatorial Auctions with Submodular Bidders. In *Proceedings of the Sixtieth Annual IEEE Foundations of Computer Science (FOCS)*.

Badanidiyuru, A.; Kleinberg, R.; and Singer, Y. 2012. Learning on a budget: posted price mechanisms for online procurement. In *Proceedings of the 13th ACM conference on electronic commerce*, 128–145.

Bade, S.; and Gonczarowski, Y. A. 2017. Gibbard-Satterthwaite Success Stories and Obvious Strategyproofness. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, EC '17, 565. New York, NY, USA: Association for Computing Machinery. ISBN 9781450345279.

Balkanski, E.; Garimidi, P.; Gkatzelis, V.; Schoepflin, D.; and Tan, X. 2022. Deterministic Budget-Feasible Clock Auctions. In Naor, J. S.; and Buchbinder, N., eds., *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9 - 12, 2022*, 2940–2963. SIAM.

Bei, X.; Chen, N.; Gravin, N.; and Lu, P. 2017. Worst-case mechanism design via bayesian analysis. *SIAM Journal on Computing*, 46(4): 1428–1448.

Budish, E.; Che, Y.-K.; Kojima, F.; and Milgrom, P. 2013. Designing random allocation mechanisms: Theory and applications. *American economic review*, 103(2): 585–623.

Christodoulou, G.; Gkatzelis, V.; and Schoepflin, D. 2022. Optimal Deterministic Clock Auctions and Beyond. In Braverman, M., ed., *13th Innovations in Theoretical Computer Science Conference, ITCS 2022, January 31 - February 3, 2022, Berkeley, CA, USA*, volume 215 of *LIPIcs*, 49:1–49:23. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

Dobzinski, S. 2007. Two randomized mechanisms for combinatorial auctions. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, 89–103. Springer.

Dobzinski, S.; Nisan, N.; and Schapira, M. 2012. Truthful randomized mechanisms for combinatorial auctions. *Journal of Computer and System Sciences*, 78(1): 15–25.

Dobzinski, S.; Ron, S.; and Vondrák, J. 2022. On the Hardness of Dominant Strategy Mechanism Design. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2022, 690–703. New York, NY, USA: Association for Computing Machinery. ISBN 9781450392648.

Dütting, P.; Gkatzelis, V.; and Roughgarden, T. 2014. The Performance of Deferred-Acceptance Auctions. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, EC '14, 187–204. New York, NY, USA: Association for Computing Machinery. ISBN 9781450325653.

Feige, U.; Flaxman, A.; Hartline, J. D.; and Kleinberg, R. 2005. On the competitive ratio of the random sampling auction. In *International Workshop on Internet and Network Economics*, 878–886. Springer.

Feldman, M.; Gkatzelis, V.; Gravin, N.; and Schoepflin, D. 2022. Bayesian and Randomized Clock Auctions. In Pennock, D. M.; Segal, I.; and Seuken, S., eds., *EC '22: The 23rd ACM Conference on Economics and Computation, Boulder, CO, USA, July 11 - 15, 2022*, 820–845. ACM.

Ferraioli, D.; Meier, A.; Penna, P.; and Ventre, C. 2019. Obviously strategyproof mechanisms for machine scheduling. In *27th Annual European Symposium on Algorithms (ESA 2019)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik.

Ferraioli, D.; Penna, P.; and Ventre, C. 2021. Two-Way Greedy: Algorithms for Imperfect Rationality. In Feldman, M.; Fu, H.; and Talgam-Cohen, I., eds., *Web and Internet Economics - 17th International Conference, WINE 2021, Potsdam, Germany, December 14-17, 2021, Proceedings*, volume 13112 of *Lecture Notes in Computer Science*, 3–21. Springer.

Ferraioli, D.; and Ventre, C. 2023a. On the Connection between Greedy Algorithms and Imperfect Rationality. In *Proceedings of the 24th ACM Conference on Economics and Computation*, EC '23, 657–677. New York, NY, USA: Association for Computing Machinery. ISBN 9798400701047.

Ferraioli, D.; and Ventre, C. 2023b. On the Connection between Greedy Algorithms and Imperfect Rationality. In Leyton-Brown, K.; Hartline, J. D.; and Samuelson, L., eds., *Proceedings of the 24th ACM Conference on Economics and Computation, EC 2023, London, United Kingdom, July 9-12, 2023*, 657–677. ACM.

Gkatzelis, V.; Markakis, E.; and Roughgarden, T. 2017. Deferred-Acceptance Auctions for Multiple Levels of Service. In *Proceedings of the 2017 ACM Conference on*

Economics and Computation, EC '17, 21–38. New York, NY, USA: Association for Computing Machinery. ISBN 9781450345279.

Goldberg, A. V.; Hartline, J. D.; Karlin, A. R.; Saks, M.; and Wright, A. 2006. Competitive auctions. *Games and Economic Behavior*, 55(2): 242–269.

Kagel, J. H.; Harstad, R. M.; and Levin, D. 1987. Information impact and allocation rules in auctions with affiliated private values: A laboratory study. *Econometrica: Journal of the Econometric Society*, 1275–1304.

Li, S. 2017. Obviously Strategy-Proof Mechanisms. *American Economic Review*, 107(11): 3257–87.

Mackenzie, A. 2020. A revelation principle for obviously strategy-proof implementation. *Games and Economic Behavior*, 124: 512–533.

Nagel, L.; and Saitto, R. 2023. A Measure of Complexity for Strategy-Proof Mechanisms. In *EC*, 1017.

Pycia, M.; and Troyan, P. 2023a. Obviously Strategyproof Mechanisms in General Environments.

Pycia, M.; and Troyan, P. 2023b. A theory of simplicity in games and mechanism design. *Econometrica*, 91(4): 1495–1526.

Qiu, F.; and Weinberg, S. M. 2024. Settling the Communication Complexity of VCG-Based Mechanisms for All Approximation Guarantees. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, STOC 2024, 1192–1203. New York, NY, USA: Association for Computing Machinery. ISBN 9798400703836.

Reiffenhauser, R. 2019. An optimal truthful mechanism for the online weighted bipartite matching problem. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, 1982–1993. SIAM.

Ron, S. 2024. Impossibilities for Obviously Strategy-Proof Mechanisms. In Woodruff, D. P., ed., *Proceedings of the 2024 ACM-SIAM Symposium on Discrete Algorithms, SODA 2024, Alexandria, VA, USA, January 7-10, 2024*, 19–40. SIAM.

Thomas, C. 2021. Classification of Priorities Such That Deferred Acceptance is OSP Implementable. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, EC '21, 860. New York, NY, USA: Association for Computing Machinery. ISBN 9781450385541.

Troyan, P. 2019. Obviously Strategy-Proof Implementation Of Top Trading Cycles. *International Economic Review*, 60(3): 1249–1261.

Vickrey, W. 1961. Counterspeculation, Auctions and Competitive Sealed Tenders. *Journal of Finance*, 8–37.