Deferred-Acceptance Heuristic Auctions*

Paul Milgrom† Ilya Segal

August 20, 2013

Abstract

We study a class of "deferred-acceptance" heuristic auctions for settings with single-minded bidders. The auctions determine allocations iteratively by rejecting the least attractive remaining bids. Strategy-proofness is achieved by letting winners' payments be their "threshold prices." We show that any deferred-acceptance heuristic auction with threshold pricing: (1) is equivalent to a clock auction in which bidders who haven't quit are accepted at their final clock prices; (2) is (weakly) group strategy-proof, as are the corresponding clock auctions for any information disclosure policies; (3) predict the same prices and assignments as the complete-information undominated Nash equilibrium of their paid-as-bid counterparts using the same heuristic allocation rule. For non-bossy deferred-acceptance heuristics, there is a unique undominated Nash equilibrium of the paid-as-bid auction, which is also the unique outcome surviving iterated deletion of weakly dominated strategies. In contrast, auctions based on optimization or greedy-acceptance heuristics generally fail properties (1)-(3).

*Segal gratefully acknowledges the support of the Toulouse Network for Information Technology.

†Department of Economics, Stanford University, Stanford, CA 94305. The authors are consultants to the US Federal Communications Commission concerning radio spectrum auction design. This research analyzes some properties of proposals to which the authors have contributed.

1 Introduction

We study a class of auctions for computationally challenging resource allocation problems for which it may be impossible to compute exactly optimal allocations. One application is to combinatorial auctions. A more recent application – and the one that motivates this study – arises from the US government's effort to reallocate frequencies currently allocated for television broadcasting to use instead for wireless broadband services. This reallocation involves purchasing television broadcast rights from some TV stations and reassigning the remaining over-the-air broadcasters to a smaller set of channels. The reassignment must be done so that no two broadcast stations are assigned to channels in ways that create interference between them. Even finding out whether a given set of broadcasters can be feasibly assigned to channels is a computationally challenging problem, equivalent to the NP-hard "graph coloring problem" (see Aardal et al. (2007) for a survey of computational approaches to this problem). The problem of selecting a feasible set of assigned stations to maximize the total broadcast value is an even harder computational problem, which cannot be solved exactly in reasonable time using today's state-of the art algorithms and hardware.

These kinds of computationally challenging mechanism design problems have been studied in the growing field of "Algorithmic Mechanism Design" (which originated with the work of Nisan and Ronen (1999)).[1] The leading approach, pioneered by Lehmann, O'Callaghan and Shoham (2002) (hereafter LOS), is to design a strategy-proof mechanism using a "greedy acceptance" heuristic algorithm.[2] For the problem of selling a set of heterogeneous goods, a "greedy acceptance" heuristic prioritizes the bids according to some "score," and iteratively accepts the highest-scoring bid that is still feasible. LOS introduce the important concept of a single-minded bidder – one who is interested in buying just one particular package of goods – and determine the unique payments to winning bidders that make truthful reporting a dominant strategy for all single-minded bidders.

[1] Economists have long been concerned about the computational properties of economic allocation mechanisms, starting at least with Hayek (1945). However, the economic literature has focused on formal modeling of communication costs (e.g. Hurwicz 1977, Mount and Reiter 1974, Nisan and Segal 2006, Segal 2007), which are trivial in the present setting of single-minded bidders, while the computational burden could be overwhelming.

[2] A number of variants of greedy-acceptance heuristic auctions has been studied – see, e.g., Mu'alem and Nisan (2008), Babaioff and Blumrosen (2008), and references therein.

In this paper, we investigate auctions based on a different class of heuristics. Like the greedy acceptance heuristics, the alternative heuristics prioritize bids for consideration, but the processing begins with the "least attractive" bids for the auctioneer instead of the "most attractive" ones. Instead of greedily accepting the most attractive bids, the alternative heuristics greedily reject the least attractive bids and, when the algorithm terminates, the bids that were not rejected are finally accepted. To create a strategy-proof auction when bidders are single-minded, each winner's payment is set to its "threshold price," which is its least attractive bid that would have still won, given the bids of the others. We call the combination a "deferred-acceptance threshold auctions." The algorithm is related to the Gale-Shapley deferred acceptance algorithm, from which we borrow the name.[3] In our setting with monetary bids, deferred-acceptance threshold auctions are also closely related to clock auctions, which offer rejected bidders an opportunity to "improve" their price offers, and at the end accept all the standing bids that are not rejected. For numerous examples of deferred-acceptance heuristics and/or clock auctions, see Kelso and Crawford (1982), Moulin (1999), Ausubel (2004), Gul and Stachetti (2000), Milgrom (2000), Hatfield and Milgrom (2005), Juarez (2009), Mehta et al. (2009), de Vries et al. (2007), Bikhchandani et al. (2011), and Ensthaler and Giebe (2009, 2010).

Despite the obvious similarity between the greedy-acceptance and deferred-acceptance algorithms – one greedily accepts attractive bids, rejecting the remainder, and the other greedily rejects unattractive bids, accepting the remainder – the classes of auctions based on these algorithms have some very different properties.[4] All deferred-acceptance threshold auctions are group

[3]In the Gale-Shapley two-sided matching algorithm, the side that receives offers rejects all but the best offers at each round and only at the end accepts the offers that were not rejected. As we show below, the deferred acceptance threshold auction can be implemented as a clock auction in which the offering side first offers its most preferred outcome and then moves down its list to less-preferred ones, with the auctioneer rejecting inferior offers along the way.

Just as we contrast the deferred acceptance procedure with greedy acceptance, the literature on matching students to schools often contrasts the Gale-Shapley deferred acceptance procedure a greedy acceptance algorithm known as the "Boston mechanism," in which the schools receiving offers immediately accept the best ones until all their slots are filled. For example, see Abdulkadiroglu, Pathak, Roth and Sonmez (2005).

[4]There is a small number of allocation rules, including the rule which takes the k highest value objects from a set of n objects ($n > k$), that can be implemented by either a greedy-acceptance or a deferred-acceptance algorithm. Obviously, the auction properties

strategy-proof and can be implemented using clock auctions, but the class of greedy-acceptance threshold auctions has neither property. All paid-as-bid auctions based on "non-bossy" deferred-acceptance heuristics are dominance solvable, but no such property holds for the class of non-bossy greedy-acceptance heuristics.[5] Another difference emerges in an outcome-equivalence property: when the same non-bossy deferred-acceptance heuristic is used to select winners for either threshold auctions or paid-as-bid auctions, the complete-information undominated Nash equilibrium outcomes for the two mechanisms coincide, but the same is not true when a greedy-acceptance heuristic is used.

In another departure from previous studies, we consider an unusually wide class of heuristics in which the prioritization of bids is not fixed in advance but can be adjusted during the auction depending on the identities of the previously rejected bidders as well as their monetary bids. This dependence is critically important for both the theoretical and practical parts of our analysis. For the theory, we use it to prove that the class of deferred-acceptance heuristic auction algorithms is exactly equivalent to the class of clock auctions (in which price adjustments may be determined by the history of bidder behavior), and that this wide class of auctions can implement any monotonic allocation rule that satisfies a "no-disposal" property and in which bids are price-theoretic substitutes. For the practice, it expands the class to include auctions in which the auctioneer faces a budget constraint (or revenue target), and also auctions that use "yardstick competition" among bidders to reduce costs (or increase revenues), as in Segal (2003). It also allows us to treat complex constraints such as those found in the spectrum reallocation problem as a routine part of the determination of bid priorities.

For concreteness, we focus on procurement ("reverse") auctions, in which the deferred-acceptance heuristic begins with an excess-supply situation and proceeds iteratively, rejecting the highest-scoring bids (with scores increasing in the bid amounts) until excess supply is eliminated. Our results are easily adapted to apply to selling ("forward") auctions, which begin with excess-demand situations and iteratively reject the lowest-scoring bids until excess demand is eliminated. This connects our analysis to all the studies of clock

coincide in those cases.

[5]Recall that an allocation rule is "non-bossy" if no change in a participant's report can alter another participant's winning/losing status without altering its own status as well. In particular, our results require that changes to *losing* bids so that they remain losing never affects which other bidders are winning.

auctions cited above, some of which focus on forward auctions.

The paper is organized as follows: Section 2 describes the general class of deferred acceptance heuristics for processing bids and gives a number of examples. Section 3 defines clock auctions in which winners are paid their final clock prices, and shows that, when bidders are restricted to use only cutoff strategies, every clock auction is equivalent to some sealed-bid deferred-acceptance threshold auction, and vice versa. Section 4 shows that, in private-value environments, clock auctions are strategy-proof and (weakly) group strategy-proof: there is no individual or coalitional deviation from truthful bidding that makes all the deviators strictly better off. This result holds regardless of what information is disclosed during the auction, from full disclosure in one extreme to no disclosure in the other. By the previous equivalence, it then follows that deferred-acceptance threshold auctions are group strategy-proof. Section 5 answers a question about the kinds of allocation rules that are exactly implementable using a deferred acceptance procedure. It shows that any monotonic allocation rule that treats bids as substitutes and has a "no-disposal" property can be implemented by some deferred-acceptance heuristic.

Section 6 fixes any deferred-acceptance heuristic and compares the performance of two auction designs – paid-as-bid and threshold auctions – under the assumption of complete information. This comparison provides one indicator of the possible cost of replacing the common paid-as-bid auction by a similar auction with dominant strategy incentives. It finds that the paid-as-bid auction has a Nash equilibrium with the same outcome as the dominant strategy solution of the threshold auction. In this equilibrium, winners bid their threshold prices – which are the same amounts they would pay in the corresponding threshold auction – and losers bid the lowest amounts exceeding their values. While a paid-as-bid auction will typically have other Nash equilibria as well, we show that for any non-bossy deferred acceptance heuristic, the paid-as-bid auction is generically dominance-solvable (i.e., its payoffs are completely determined by iterated elimination of dominated strategies) and that the dominance solution coincides with the equilibrium described above.[6] Moreover, in that case, the same outcome (allocation and prices)

[6]Bids made by losing bidders are often not uniquely determined just by iterated elimination of dominated strategies. When the allocation rule is non-bossy, the indeterminacy of losing bids does not affect the auction outcome, but when the allocation rule is bossy, the winners' identities can depend on the losing bids and iterated dominance may fail to determine a unique auction outcome.

obtains in every Nash equilibrium in undominated strategies.[7] The distinctiveness of the dominance solvability finding is highlighted by a partial converse: any dominance-solvable paid-as-bid auction that selects winners using a monotonic, non-bossy allocation rule must implement the same outcome mapping as some deferred-acceptance heuristic.

In Section 7, we demonstrate by means of examples that the properties of group strategy-proofness and paid-as-bid outcome equivalence do not hold generally when winner selection is by either a greedy-acceptance heuristic or optimization.

2 Heuristic Sealed-Bid Auction

Let N be the set of bidders. In the auction, each bidder either "wins" (which means that his bid to supply a given good or set of goods is "accepted") or "loses" (which means that his bid is rejected). We restrict attention to auctions in which winners receive payments but losers do not. The preferences of each bidder i depend on whether he wins or loses, and, when he wins, on the payment p_i. We assume that these preferences are strictly increasing in the payment, and there exists some payment v_i that makes him indifferent between winning and losing, which we call his "value". (For unmixed outcomes, such a preference can be expressed by a quasilinear utility p_i when the bidder wins and v_i when he loses.) The set of bidder i's possible values is $[0, \overline{v}_i]$ for each i.

A mechanism requests each bidder i to submit a bid $b_i \in B_i \subseteq \mathbb{R}$ and generates a set of $A \subseteq N$ of accepted bids for every bid profile $b \in B = \Pi_i B_i$. We restrict each bid space B_i to be a closed set such that $\sup B_i > \overline{v}_i$.[8]

[7] These outcome equivalence results are related to the findings of Bernheim and Whinston (1986) about equivalence between optimizing paid-as-bid auctions and Vickrey auctions for the case in which bidders are substitutes. One connection is that when bidders are substitutes, the Vickrey auction is implementable by a deferred acceptance heuristic and our findings apply to general deferred-acceptance heuristics. To isolate a particular Nash equilibrium outcome, Bernheim and Whinston (1986) refine equilibrium using coalition-proofness while we instead use iterated deletion of dominated strategies.

[8] The last restriction ensures that each bidder prefer to participate in the auction. Alternatively, we could restrict attention to auctions in which each bidder has a "non-participation bid" that loses against any profile of other bids. In this case, a bidder i with value $v_i \geq \sup B_i$ will submit a non-participation bid. The bidder's maximum bid that has the possibility of winning could then be interpreted as that bidder's "reserve price".

(Below, we often further restrict B_i to be finite.) Let $\alpha : B \to 2^N$ denote the winner determination rule or "allocation rule" of the mechanism: $\alpha(b) \subseteq N$ is the set of winners generated for bid profile $b \in B$.

A deferred-acceptance heuristic is a particular kind of mechanism described by a set of scoring functions, as follows. For each set $A \subseteq N$ of active bidders and each bidder $i \in A$, there is a *scoring function* $s_i^A : B_i \times B_{N \setminus A} \to \mathbb{R}_+^A$ that is nondecreasing in its first argument. The heuristic then operates as follows. Let $A_t \subseteq N$ denote the set of active bids in stage t. Initialize $A_1 = N$. For each $t \geq 1$, if $s_i^{A_t}\left(b_i, b_{N \setminus A_t}\right) = 0$ for all $i \in A_t$ then stop and output A_t, otherwise let $A_{t+1} = A_t \setminus \arg\max_{i \in A_t} s_i^{A_t}\left(b_i, b_{N \setminus A_t}\right)$ and continue. Intuitively, the heuristic process is one of iteratively deleting the least desirable (highest scoring) bids until only zero scores remain.

2.1 Payments and Strategy-proofness

A *payment rule* is a function $p : B \to \mathbb{R}^N$ specifying that agent i receives $p_i(b)$. A *sealed-bid auction* is a triple $\langle B, \alpha, p \rangle$ such that $p_i(b) = 0$ whenever $i \in N \setminus \alpha(b)$ (meaning that losing bidders are not paid).

Since we will often use finite bid sets, it is convenient to replace the usual notion of "truthful" bidding by a concept of "strategy-proofness" that applies even when some possible bidder values do not correspond to feasible bids. According to our definition, an auction is strategy-proof if it is always optimal for a bidder to round up its value to next lowest allowable bid. With this definition, if the sets of possible values and bids are both the same interval of real numbers, then strategy-proofness and truthfulness coincide.

Definition 1 *The sealed-bid auction $\langle B, \alpha, p \rangle$ is* strategy-proof *if for every bidder i, $v_i \in [0, \bar{v}_i]$, and $b_{-i} \in B_{-i}$, it is optimal for bidder i to bid $v_i^+ \equiv \min\{b_i \in B_i : b_i > v_i\}$ (and in particular the minimum must exist).*

Definition 2 *The allocation rule α is* monotonic *if and only if $i \in \alpha(b_i, b_{-i})$ and $b_i' < b_i$ imply $i \in \alpha(b_i', b_{-i})$.*

With these definitions, a standard argument implies the following:

Lemma 3 *A sealed-bid auction $\langle B, \alpha, p \rangle$ is strategy-proof if and only if α is monotonic and payments satisfy the following formula for all $b \in B$, $i \in \alpha(b)$:*

$$p_i\left(b_{-i}\right)=\sup\left\{b_i'\in B_i: i\in\alpha\left(b_i',b_{-i}\right)\right\}. \tag{1}$$

It is easy to see that any deferred-acceptance heuristic generates a monotonic allocation rule and that the corresponding threshold prices (1) to the winners can be computed as follows: Start with $p_i^0=\sup B_i$ for all i, and then for each round $t\geq 1$, compute

$$p_i^t(b)=\min\left\{p_i^{t-1},\sup\left\{b_i'\in B_i: s_i^{A_t}\left(b_i',b_{N\backslash A_t}\right)<s_j^{A_t}\left(b_j,b_{N\backslash A_t}\right)\text{ for } j\in A_t\backslash A_{t+1}\right\}\right\}$$

for every bidder $i\in A_{t+1}$. In the final round of the algorithm, for every winner $i\in A^T$, $p_i^T(b)$ is the winner's threshold price.[9]

Inspection of the formula shows a consequential property of the threshold prices for deferred-acceptance heuristics: holding fixed the final set of winners A_T, winning bidders' threshold prices depend on the losing bids $b_{N\backslash A_T}$ but *not* on the winning bids b_{A_T}. It follows that no winning bidder can affect another winner's threshold price except by changing to a losing bid.

2.2 Examples

Example 4 (Feasibility Constraint) *Let $F\subseteq 2^N$ denote the set of sets of bidders that could be feasibly accepted, and assume that $N\in F$, so that the procurement goal is achievable. To ensure that the heuristics maintains feasibility, we require that $s_i^A\left(b_i,b_{N\backslash A}\right)>0$ only if $A\backslash\{i\}\in F$, and also that there are no ties, i.e., $s_i^A\left(b_i,b_{N\backslash A}\right)\neq s_j^A\left(b_j,b_{N\backslash A}\right)$ for all $i\neq j$, A, $b_i,b_j,b_{N\backslash A}$.*

We say that the heuristic has perfect feasibility checking *if $s_i^A\left(b_i,b_{N\backslash A}\right)>0$ if and only if $A\backslash\{i\}\in F$ – i.e., it stops only when all active bids are infeasible to reject. In some settings, however, perfect feasibility checking may be too computationally challenging, and imperfect checking must be used instead. For example, in the FCC's spectrum-clearing problem, to check whether a given set A of bidders is in F requires checking whether there exists an assignment of the rejected bidders $N\backslash A$ to available channels that satisfies all interference constraints, and this is an NP-hard problem. When a feasibility checker has a limited time to run, it may generate three possible outputs: (i)*

[9] These round-by-round computations can be integrated with the heuristic in one single calculation step.

establish that $A \backslash \{i\} \in F$ by generating a feasible assignment of the rejected bidders to channels, (ii) establish that $A \backslash \{i\} \notin F$ by generating a proof that such a feasible assignment does not exist, and (iii) be timed out before generating either (i) or (ii). In case (i), we set $s_i^A \left(b_i, b_{N \backslash A}\right) > 0$, while in cases (ii) and (iii), we need to set $s_i^A \left(b_i, b_{N \backslash A}\right) = 0$, to guarantee that the heuristic yields a feasible assignment.

The next two examples show two reasons why it may be useful to condition the scoring functions on the rejected bids. The first reason is to incorporate a budget restriction that makes the auctioneer reject additional bids when the current overall costs are too high. The second reason is to create "yardstick competition" among bidders by inferring reasonable reserve prices for the active bidders from the rejected bids.

Example 5 (Budget or Payment Constraint) *Suppose that the total payment to the winners cannot exceed $R(A)$ when the set of accepted bids is A. For example, in the FCC case, payments to broadcasters are limited by the revenue obtained from selling the cleared spectrum in the forward auction, net of some expenditures required by statute or regulations. Since the FCC may be initially uncertain about how much spectrum it can clear subject to a net payment constraint, it might set a sequence of n possible procurement goals represented by feasible sets: $F_1 \subset ... \subset F_n$ with the corresponding forward auction revenues $R_1, ..., R_n$, so that the maximum forward auction revenue achieved by accepting set A of bids is $R(A) = \max\{R_k : 1 \leq k \leq n,\ A \in F_k\}$.) Then in every round t, the algorithm could look at the total threshold prices that would have to be paid to each of the active bidders in A_t if the algorithm were to stop in this round, and continue to the next round if this total exceeds $R(A_t)$. For a simple example, if scores are based on functions $\sigma_i(b_i) > 0$ (independent of the comparison set A and of others' bids), the threshold price that would have to be paid to a currently active bidder $i \in A$ if the heuristic were to stop right away can be calculated as:*

$$p_i\left(b_{N \backslash A}\right) = \sup\left\{b_i' \in B_i : \sigma_i\left(b_i'\right) < \max_{j \in N \backslash A} \sigma_j\left(b_j\right)\right\}.$$

To add a "stopping rule" that allows the auction to end only if the budget constraint is met, we let for each $i \in A$,

$$s_i^A\left(b_i, b_{N \backslash A}\right) = \begin{cases} 0 & \text{if } \sum_{j \in A} p_j\left(b_{N \backslash A}\right) \leq R(A), \\ \sigma_i(b_i) & \text{otherwise.} \end{cases}$$

This can be viewed as a "revenue-sharing"' problem, which is the mirror image of the "cost-sharing"' problem of Moulin (1999) and Mehta et al. (2009). Our formulation permits the auction to generate revenue in excess of $R(A)$ *to be absorbed by the auctioneer, but it is possible to modify it to require revenue to be exactly* $R(A)$. *(This possibility will become clearer once we introduce clock auctions and show (Proposition 8) that for every clock auction there exists an equivalent deferred-acceptance threshold auction.)*

Example 6 (Reference Pricing) *Suppose the auctioneer cares about the expected total profit, and that his gross profit for acquiring each bidder is* π. *Suppose that bidders' values are drawn i.i.d. from a distribution that is unknown to the auctioneer. In this symmetric setting we can focus without loss on symmetric auctions. If we consider symmetric deferred-acceptance heuristics that do not condition on the rejected bids, these heuristics accept all bids above some fixed reserve price* p^* *(e.g. setting* $s_i^A(b_i) = \max\{v_i - p^*, 0\}$*), which yields an expected profit on each bidder* i *of* $(\pi - p^*)\Pr\{v_i \leq p^*\}$, *where the probability is calculated based on the auctioneer's prior. Note, however, that these expected profits could be improved by conditioning the reserve price* $p_A^*\left(b_{N\setminus A}\right)$ *in each round on the rejected bids (i.e., letting* $s_i^A\left(b_i, b_{N\setminus A}\right) = \max\left\{v_i - p_A^*\left(b_{N\setminus A}\right), 0\right\}$*). For example, the reserve price could be the optimal price for the beliefs about the distribution of values of the active bidders that are updated based on* $b_{N\setminus A}$. *(Note: the expected profit-maximizing threshold auction, described in Segal (2003), implements the allocation rule* $\alpha(b) = \{i \in N : b_i \leq p(b_{-i})\}$ *where* $p(b_{-i}) \in \arg\max_p (\pi - p)\Pr\{v_i \leq p | b_{-i}\}$. *If the family of possible distributions of bidders' values is ordered in the likelihood ratio ordering, then* $p(b_{-i})$ *is nondecreasing in* b_{-i}, *hence allocation rule* α *has the substitute property. If in addition the range of* α *has no disposal, then by Proposition 14 below* α *can be implemented as a heuristic auction.)*

3 Clock Auction

Informally, a ("descending") clock auction is a dynamic mechanism that proposes a declining sequence of price offers to each bidder, with each offer followed by a decision period in which the bidder whose price has been strictly reduced can decide to exit or continue. Bidders have never exited are called "active"; others are called "inactive." Bidders who remain active when their prices are reduced are said to "accept" the lower price. When the auction ends, the active bidders become the winners and they are paid their last (lowest) accepted prices. Different clock auctions are distinguished by their pricing functions, which determines the sequence of prices to offer the several bidders.

We formalize this mechanism as follows: A period-t history consists of the sets of active bidders in all periods up to period t – i.e., $A^t = (A_1, ..., A_t) \in (2^N)^t$ such that $A_t \subseteq ... \subseteq A_1 = N$. Let H denote the set of all such histories. A descending clock auction is defined by a price mapping $p : H \to \mathbb{R}^N$ such that for all $t \geq 2$ and all A^t, $p(A^t) \leq p(A^{t-1})$. (Note that we reuse the p notation here to represent the pricing in the clock auction: it had earlier referred to pricing in the threshold auction.)

The clock auction initializes $A_1 = N$. In each period $t \geq 1$, given history A^t, it offers prices $p(A^t)$ to bidders. If $p(A^t) = p(A^{t-1})$, the auction stops and bidder i is a winner if and only if $i \in A_t$, and in that case i is paid $p_i(A^t)$. If $p(A^t) \neq p(A^{t-1})$, then every bidder in A_t chooses whether to exit the set of active bidders.[10] Letting $E_t \subseteq A_t$ denote the set of bidders who choose to exit, the auction continues in period $t+1$ with the new set of active bidders $A_{t+1} = A_t \backslash E_t$ and the new history $A^{t+1} = (A^t, A_{t+1})$. We say that the clock auction is *finite* if there exists some T such the auction always stops by period T.

To complete the description of the auction as an extensive-form mechanism, we also need to describe bidders' information sets. We allow general information disclosure: bidder i observes some signal $\sigma_i(A^t)$ in addition to his current price $p_i(A^t)$ in history A^t.

A strategy for bidder i in a clock auction is a *cutoff strategy* with cutoff b_i if it specifies exit if and only if $p_i(A^t) < b_i$, for some $b_i \leq p_i(N)$. Note in particular that every cutoff strategy accepts the opening price. The next two results show that clock auctions in which bidders are restricted to cutoff strategies (for example, in which they must use proxy bidders with cutoff strategies) are equivalent to deferred-acceptance threshold auctions, meaning that the mapping from bid/cutoff profiles to allocations and prices are the same for both auctions.[11] We restrict attention to heuristics with finite bid

[10] In a variant of the auction, only active bidders whose prices are decremented in the current round may be permitted to exit. Although the results below are the same for the auction in the main text and this variant, there is one important difference: with a feasibility constraint, this variant ensures that the clock auction yields a feasible outcome in cases like Example 4.

[11] If we were to implement a deferred-acceptance threshold auction as a multi-round

spaces and to finite clock auctions.[12]

Proposition 7 *For every deferred-acceptance heuristic with finite bid spaces and threshold pricing, there exists an equivalent finite clock auction in which bidders are restricted to cutoff strategies.*

Proof. Given bid spaces $B_1, .., B_N$, for each $v \in \mathbb{R}$, let and $v^+ = \min\{b_i \in B_i : b_i > v\}$ and $v^- = \begin{cases} \max\{b_i \in B_i : b_i < v\} & \text{if } v_i > \min B_i, \\ \min B_i - 1 & \text{otherwise.} \end{cases}$. Let the opening prices be $p_i(N) = \max B_i$ for each i. Given a deferred-acceptance heuristic auction with scoring rule s, we construct an equivalent clock auction as follows: The price reduction rule in the clock auction reduces the price to every highest-scoring active bidder by the minimal amount, while leaving prices unchanged for the other bidders:

$$\begin{aligned} p_i\left(A^t\right) &= p_i\left(A^{t-1}\right)^- \text{ if } i \in \arg\max_{j \in A_t} s_j^{A_t}\left(p_j\left(A^{t-1}\right), p_{N \setminus A_t}\left(A^t\right)^+\right) \\ p_i\left(A^t\right) &= p_i\left(A^{t-1}\right) \text{ otherwise.} \end{aligned}$$

Note in particular that the auction maintains $p_i(A^t) = p_i(A^{t-1})$ for all $i \in N \setminus A_t$ – thus memorizing the prices rejected by bidders who have quit, so that their cutoffs can be inferred as $p_i(A^t)^+$.

Then equivalence is easy to see: First, for every history of the clock auction, the next set of bidders to quit in the clock auction is the set of bidders who have the maximum scores among the set of active bidders, so the set of winners is the same in both auctions. Second, if any winning bidder had said "no" to any higher price, it would have exited, so each bidder's final clock price is the highest cutoff it could use to be winning – its threshold price.

Proposition 8 *For every finite clock auction in which bidders are restricted to cutoff strategies, there exists an equivalent deferred-acceptance heuristic with finite bid spaces and threshold prices.*

procedure in which some information is disclosed to active bidders between rounds and they are allowed to improve their bids, then the resulting mechanisms would be "survival auctions" like those proposed by Fujishima et al. (1999) for more specific settings. These auctions are strategically equivalent to clock auctions without the restriction to cutoff strategies.

[12] This restriction avoids technical difficulties associated with describing continuous time auctions and with defining dominance solvability for infinite games.

■

Proof. Given a finite clock auction p, we construct bid spaces and a scoring rule to create an equivalent deferred-acceptance heuristic. We take each bidder i's bid space to be $B_i = \{p_i(h) : h \in H\}$ – the set of possible prices agent i could face in the clock auction (which is a finite set in a finite clock auction).

Next, we construct the scoring rule in the following manner: Holding fixed a set of bidders $S \subseteq N$ and their bids $b_S \subseteq \mathbb{R}^S$, let $A_t(S, b_S)$ denote the set of active bidders in the clock auction at round t in which every bidder $j \in S$ uses cutoff strategy b_j and every bidder from $N \backslash S$ never exits. Formally, initialize $A_1(S, b_S) = N$ and iterate by setting

$$A_{t+1}(S, b_S) = A_t(S, b_S) \setminus \left\{ j \in S : b_j > p_j\left(A^t(S, b_S)\right) \right\}.$$

This gives an infinite sequence $\{A_t(S, b_S)\}_{t=1}^{\infty}$, but the sets start repeating at some point (when the clock auction stops).

Now for given A, $b_{N\backslash A}$, $i \in A$, and b_i, define the score of agent i as the inverse of how long he would remain active in clock auction if he uses cutoff b_i and all bidders from $N\backslash A$ use cutoffs $b_{N\backslash A}$, while bidders in $A\backslash\{i\}$ never quit:

$$s_i^A\left(b_i, b_{N\backslash A}\right) = 1/\sup\left\{t \geq 1 : i \in A_t\left(\{i\} \cup (N\backslash A), \left(b_i, b_{N\backslash A}\right)\right)\right\}.$$

(Note that the score is $1/\infty = 0$ in cases in which the auction stops with agent i still active.) This score is by construction nondecreasing in b_i. Also by construction, given a set A of active bidders, the set of bidders to be rejected by the heuristic in the next round ($\arg\max_{i \in A} s_i^A\left(b_i, b_{N\backslash A}\right)$) is the set of bidders who would quit the soonest in the clock auction given that the inactive bidders have used cutoffs $b_{N\backslash A}$. If no more bidders would exit the auction, then all active bidders have the score of zero, so the heuristic stops. Finally, as argued above, the winners' clock auction prices are their threshold prices: the winner would have lost by using any higher cutoff in B_i than its clock auction price. ■

4 (Group) Strategy-proofness

Definition 9 *In a clock auction, agent i with value v_i is said to 'bid truthfully" if he accepts clock price if and only if $p_i(h) > v_i$. (Equivalently, if the agent uses a cutoff strategy with cutoff $v_i^+ = \min\{p_i(h) : h \in H,\ p_i(h) > v\}$.)*

Definition 10 *An auction is "weakly group strategy-proof" if for every profile of values v and every set of players $S \subseteq N$ and every strategy profile σ_S of these players, at least one bidder in S has a weakly higher payoff from the profile of truthful bids v_N than from the strategy profile $(v^+_{N\setminus S}, \sigma_S)$.*

Remark 11 *Clock and threshold auctions are not generally "strongly" group strategy-proof, because a bid increase by a losing bidder that increases a winner's threshold price is strictly profitable for the winner and weakly profitable for the loser.*

Clock auctions can have various information disclosure policies, leading to a potentially large set of strategies for bidders, but always including the cutoff strategies. The definition of group strategy-proofness applies to all such auctions.

Proposition 12 *Every finite clock auction (with any information disclosure) is weakly group strategy-proof.*

Proof. Consider the first stage of clock auction affected by a group deviation. If at that stage, the deviation is by a bidder who chooses to exit, then his deviation payoff is zero, so he does not benefit from the group deviation. The other possibility is that the deviation is by a bidder who chooses not to exit at a price equal to or below his truthful value, but such a bidder either eventually exits or wins and receives his final clock price, which in a descending clock auction cannot be higher and so cannot exceed his value. Hence, this deviator's payoff is non-positive. In both cases, at least one participant in the group deviation fails to gain from the deviation. ■

The preceding argument is independent of the information policy in the auction, so it works for larger sets of strategies than just the cutoff strategies. In particular, the clock auction is still group strategy-proof when bidders are restricted to play cutoff strategies. Combining the two previous propositions, we get

Corollary 13 *Any deferred-acceptance auction with threshold prices and finite bid spaces is weakly group strategy-proof.*

It is possible to establish the result directly, without using the equivalence to clock auctions or restricting attention to finite bid spaces.

5 Substitutability -> Heuristic implementation

We say that a clock auction *implements allocation rule* $\alpha : B \to 2^N$ if (i) $\{p_i(h) : h \in H\} = B_i$ for each i and (ii) the auction generates $\alpha(b)$ when bidders use cutoff strategies with cutoffs $b \in B$. Also, the allocation rule α is *monotonic* if $i \in \alpha(b)$ and $b'_i < b_i$ implies $i \in \alpha(b'_i, b_{-i})$. It *has substitutes* if $i \in \alpha(b)$ and $b'_j > b_j$ for some $j \neq i$ implies $i \in \alpha\left(b_j, b'_j\right)$.

A set $S \subseteq 2^N$ of subsets of N has *no disposal* if for all $A, A' \in S$, $A \subseteq A'$ implies $A = A'$.

Proposition 14 *With finite bid spaces, any monotonic allocation rule α with substitutes whose range $\alpha(B)$ has no disposal can be implemented with a clock auction or a deferred acceptance heuristic.*

Proof. α can be implemented with a clock auction described as follows: For each i, set $p_i(N) = \max B_i$ and then in each period t, set

$$\begin{aligned} p_i\left(A^t\right) &= p_i\left(A^{t-1}\right)^- \text{ if } i \in A_t \backslash \alpha\left(p_{A_t}\left(A^{t-1}\right), p_{N\backslash A_t}\left(A^{t-1}\right)^+\right) \\ p_i\left(A^t\right) &= p_i\left(A^{t-1}\right) \text{ otherwise.} \end{aligned}$$

(That is, decrement prices to those bidders who wouldn't win given the current best offers – the current prices for the active bidders, and the last prices accepted by the bidders who have exited.)

To see that this auction implements α, observe that if bidders use cutoff strategies with cutoffs $b_i \in B_i$, then a bidder $i \in \alpha(b)$ can never exit the auction: when $i \in A_t$ and he is offered price $p_i(A^t) = b_i$, we will have $p_{A_t\backslash\{i\}}(A^{t-1}) \geq b_{A_t\backslash\{i\}}$ and $p_{N\backslash A_t}(A^{t-1})^+ = b_{N\backslash A_t}$, hence the substitute property and $i \in \alpha(b)$ imply $i \in \alpha\left(p_{A_t}(A^{t-1}), p_{N\backslash A_t}(A^{t-1})^+\right)$ and so his price is not decremented. Thus, we have $\alpha(b) \subseteq A_t$ throughout the auction. On the other hand, when the auction stops we have $A_t \subseteq \alpha\left(p_{A_t}(A^{t-1}), p_{N\backslash A_t}(A^{t-1})^+\right)$, and putting together with the previous inclusion and using the no-disposal of the range of α implies $\alpha(b) = A_t = \alpha\left(p_{A_t}(A^{t-1}), p_{N\backslash A_t}(A^{t-1})^+\right)$, hence the auction implements α. ■

The assumption of substitutes is not dispensable in the above proposition: in Example 22 below, we will see a monotonic allocation rule whose range has no disposal that cannot be implemented by a clock auction. While many deferred-acceptance heuristic allocation rules do satisfy substitutes, not all of them do. For example, consider the allocation rule $\alpha(b_1, b_2) = \begin{cases} \{1,2\} & \text{if } b_1 < 1, \\ \varnothing & \text{otherwise.} \end{cases}$ This allocation does not have substitutes, but is implementable with the deferred-acceptance heuristic with the scoring rule $s_1^{\{1,2\}}(b_1) = \max\{b_1 - 1, 0\}$, $s_2^{\{2\}}(b_2, b_1) = 1$, and $s_2^{\{1,2\}}(b_2) = s_1^{\{1\}}(b_2, b_1) = 0$.)

The no-disposal assumption is also indispensable, which is illustrated by the allocation rule $\alpha(b) = \arg\min_{i\in N} b_i$, with $B_1 = ... = B_N$ (so that ties exist, and in case of ties all the tied bidders win). Then there is no clock auction implementing α.[13] The no-disposal assumption is satisfied, e.g., in Example 4 if the feasible set F is *comprehensive* (meaning that $A \in F$ implies $A' \in F$ whenever $A \subseteq A'$) and the heuristic has perfect feasibility checking. But not all heuristic allocation rules satisfy it (e.g., in Example 5, the heuristic may accept a set A' when bids are low and a set $A \subset A'$ when bids are high).

6 Pay-as-Bid: Full-info equivalence

Recall that for any finite bid space B and allocation rule α, the threshold prices for winners are given by $p_i(b_{-i}) = \max\{b_i' \in B_i : i \in \alpha(b_i', b_{-i})\}$. In particular, $i \in \alpha(p_i(b_{-i}), b_{-i})$.

Proposition 15 *Every paid-as-bid deferred-acceptance auction with finite bid sets B_i for all values $v_i < \max B_i$ has a complete-information Nash equilibrium profile in which, for each $i \in N$, the bids are $b_i = \max\left\{v_i^+, p_i\left(v_{-i}^+\right)\right\}$ and in which the resulting allocation is $\alpha(b) = \alpha(v^+)$.*

Proof. Since changing accepted bids so that they are still accepted does not affect the deferred-acceptance heuristic's outcome, we have $A \equiv \alpha(v^+) = \alpha\left(p_A(v^+), v_{N\setminus A}^+\right) = \alpha(b)$ and $p_i(b_{-i}) = p_i\left(p_{A\setminus i}(v^+), v_{N\setminus A}^+\right) = p_i\left(v_{-i}^+\right) \geq$

[13] Indeed, any such auction would start with equal prices, and it would then not be "safe" to reduce any price: if all bidders have set their cutoff equal to the common price, then the reduction would eliminate any affected bidder. On the other hand, if some one bidder has bid below the common price, then failing to reduce his price prevents the algorithm from ever identifying that bidder.

v_i^+ for each $i \in A$. Now, we verify that the bids constitute a Nash equilibrium. Every bidder $i \in A$ is winning and receiving payment of $b_i = p_i\left(v_{-i}^+\right) = p_i\left(b_{-i}\right) \geq v_i^+$, and any larger bid by i would be losing, so a winning bidder i has no profitable deviation. Every bidder $i \in N\backslash A$ is losing with its bid of v_i^+, and so any winning bid for i earns a negative payoff: a losing bidder has no profitable deviation. ■

Next, we introduce a pair of standard definitions.

Definition 16 *An auction is* dominance-solvable *in state v if under full information, iterated deletion of (weakly) dominated strategies in any order yields a unique payoff profile.*

For "generic" values ($v_i \notin B_i$ for each i), a unique payoff profile implies a unique outcome (allocation and winning bids).

Definition 17 *Assignment rule α is "non-bossy" if for any $i \in N$, $b \in B$ and $b_i' \in B_i$, $\alpha\left(b_i', b_{-i}\right) \cap \{i\} = \alpha\left(b\right) \cap \{i\}$ implies $\alpha\left(b_i', b_{-i}\right) = \alpha\left(b\right)$.*

Non-bossiness means simply that a bidder cannot affect others' allocations without changing his own allocation. Some deferred acceptance heuristics are non-bossy, but as our "reference pricing" example illustrates, some are not. In a deferred acceptance-heuristic, a winner who changes its bid without changing its winning status (that is, agents $i \in \alpha\left(b_i, b_i\right) \cap \alpha\left(b_i', b_i\right)$) can never affect others' winning status, but because bidder's scores can depend on losing bids, a loser who changes to a different losing bid ($i \notin \alpha\left(b_i, b_i\right) \cup \alpha\left(b_i', b_i\right)$) may affect the set of winners.

Proposition 18 *Consider a paid-as-bid auction with a monotonic, non-bossy assignment rule α and finite bid spaces B. Say that a value profile v is "generic" if for each i, $v_i \in \left[0, \bar{v}_i\right] \backslash B_i$.*

(i) The auction is pure-strategy dominance-solvable for all generic value profiles if and only if α can be implemented via a deferred-acceptance heuristic.

(ii) In this case, for every generic value profile, the unique payoff profile surviving iterated deletion of dominated strategies is also the unique (pure or mixed) Nash equilibrium payoff profile in undominated strategies.

(iii) In this case, one strategy profile that survives iterated deletion of dominated strategies and is a Nash equilibrium in undominated strategies is the one described in Proposition 15.

Remark 19 *To see why we need α to be non-bossy, suppose that $N = 2$, $B_1 = \{1, 3\}$, $B_2 = \{2, 4\}$. Let $\alpha(b) = \{1\}$ if $b_2 = 4$, and $\alpha(b) = \varnothing$ otherwise. With $v_1 < 3$, bidder 1's dominant strategy is to bid 3, but bidder 2 has no dominated strategies. The strategy profiles that survive iterated elimination in this example are the same as the undominated Nash equilibria: the comprise $(3, 4)$ (in which bidder 1 wins) and $(3, 2)$ (in which there is no winner). Generally, the concepts of iterated dominance and undominated Nash equilibrium do not nail down the losing bids in these paid-as-bid auction games, which is problematic when the losing bids can influence the outcome. Like the example above, bossy rules are also used in our earlier examples of budget constraints and yardstick competition. For a unique prediction in such cases, one would need to use an equilibrium refinement to restrict the losing bids. In contrast, when α is non-bossy, one can deduce the unique auction outcome just by reasoning about winning bids, as we do in the proof below.*

Proof. For the "if" direction of (i), recall from Proposition 7 that any assignment rule α that is implementable via a deferred-acceptance heuristic is also implementable with a clock auction in which bidders use cutoff strategies with cutoffs corresponding to their bids in the deferred-acceptance heuristic. Furthermore, we can implement assignment rule α *with paid-as-bid pricing* using the following "two-phase clock auction": In phase 1, the the clock auction described above is run to determine the set of winners. In phase 2, the payments to the winners are determined by allowing prices to continue falling (through points in B_i) until all bidders "quit", with the winners being paid the last prices they accept. The two-phase clock auction game in which bidders use the cutoff profile b obviously leads to the same outcome as the paid-as-bid sealed-bid auction game based on the deferred acceptance algorithm in which the bid profile is b.

If the assignment rule is non-bossy, then for generic values $v_i \in \mathbb{R} \backslash B_i$ the game satisfies the TDI condition of Marx and Swinkels (1997), and so the payoffs profiles surviving iterated deletion do not depend on the order of deletion: hence deleting dominated or equivalent strategies in any order leads to the same set of possible outcomes.[14] We specify the following deletion process:

[14] We say that given strategy sets $\hat{B}_i \subseteq B_i$ for each i two strategies $b_i, b_i' \in \hat{B}_i$ of agent i are equivalent if $\alpha(b_i, b_{-i}) = \alpha(b_i', b_{-i})$ for all $b_{-i} \in \hat{B}_{-i}$. Obviously deletng strategies that are equivalent to surviving ones does not affect the solution to iterated deletion of dominated strategies. Note furthermore that given non-bossiness, such equivalence obtains whenever agent i's own allocation does not change, i.e. $\alpha(b_i, b_{-i}) \cap i = \alpha(b_i', b_{-i}) \cap i$ for

Begin by deleting for each agent i all the bids/cutoffs $b_i < v_i^+$ (which are either dominated by or equivalent to the bid v_i^+). In the game that remains after these initial deletions, every bidder strictly prefers any outcome in which it wins to any in which it loses. We specify the next deletions inductively by referring to the sequence of prices $\{p(A^t)\}$ that would emerge during phase 1 if each bidder were to use the cutoff strategy v_i^+. At the beginning of each step t of our iterated deletion process, the set of strategies remaining to each bidder i is $\hat{B}_i^{t-1} = B_i \cap [v_i^+, \max\left\{v_i^+, p_i\left(A^{t-1}\right)\right\}]$. As the prices are reduced to $p\left(A^t\right)$, for each bidder i all the cutoffs $b_i \in \hat{B}_i^{t-1}$ such that $b_i > \max\left\{v_i^+, p_i\left(A^t\right)\right\}$ are sure to lose and are therefore either dominated by or equivalent to the cutoff v_i^+, hence we can let $\hat{B}_i^t = B_i \cap [v_i^+, \max\left\{v_i^+, p_i\left(A^t\right)\right\}]$. The iterations continue until phase 1 ends and the winners are determined at the end of some iteration T to be $\alpha(\max \hat{B}^T)$. For each agent i, if $\hat{B}_i^T$ is not a singleton, then its largest element, $\max \hat{B}_i^T = \max(v_i^+, p_i\left(A^T\right))$, is dominant in the game with just the bids $\hat{B}^T$ (because it wins at the highest price). So, we may do one more round of deletions, taking $\hat{B}_i^{T+1} = \{\max(v_i^+, p_i\left(A^T\right))\}$. Hence, the single outcome of iterative elimination of undominated cutoffs is the one for the bid profile $(\max(v_i^+, p_i\left(A^T\right)))_{i \in N}$.

For (ii), fix an undominated mixed Nash equilibrium profile. For each bidder i with a zero equilibrium payoff, all bids of v_i^+ or more must be always losing. Hence, by non-bossiness, we may replace every such bidder i's bids by the pure strategy bid v_i^+ to obtain another mixed strategy profile σ with the same distribution of outcomes. We show below that σ is actually a pure strategy bid profile, and specifically it is the profile $(\max(v_i^+, p_i\left(A^T\right)))_{i \in N}$ that results from iterated elimination of weakly dominated strategies, as described above.

For any bidder i with strictly positive equilibrium expected payoffs, all bids in the support of σ_i have positive expected payoffs, so all must win with a positive probability against σ_{-i}. Consider the maximum bid profile in the support of σ. Referring to the clock auction process, we infer that if any positive-payoff bidder's bid is losing for that profile, then it is losing for all profiles in the support of σ, which contradicts positive expected payoffs. Since reducing a winning cutoff/bid in the clock auction does not affect the allocation, for every bid profile in the support of σ, the positive-payoff players are the winners. Since the highest always-winning bid earns strictly more than any lower winning bid, this further implies that the winners' equilibrium

all $b_{-i} \in \hat{B}_{-i}$.

mixtures are degenerate: winning bidders play pure strategies. Therefore, σ assigns probability one to some single bid profile b.

Next, we claim that the iterative deletions described in the proof of (i) above do not delete any of the component bids in b. Phase I of the iterative deletion procedure deletes only bids above v_i^+ for zero-payoff bidders and only always-losing bids for positive-payoff bidders, so all the component bids in b survive that phase. Phase II deletes all but the highest remaining bid of each winning bidder: the lower bids are never best replies to the highest surviving bids (they always win, but they are paid less). Hence, the full procedure never deletes any component bid in the profile b. It follows that $b = (\max(v_i^+, p_i\left(A^T\right)))_{i \in N}$ and that the outcome of b is the outcome of every undominated Nash equilibrium.

To prove (iii): in the surviving bid profile b, each agent $i \in A^T$ bids its threshold price, which is $p_i\left(A^T\right) \leq v_i^+$, while each $i \in N \backslash A^T$ bids v_i^+, which is by definition above its threshold price. Thus by Proposition 15 it is a Nash equilibrium and it contains only undominated strategies, and as argued above it survives iterated deletion of dominated strategies.

It remains to prove the "only if" direction of (i): we assume that the auction is dominance solvable and relate a sequence of sets $\hat{B}\left(A^t\right)$ surviving a number of rounds of iterated elimination of dominated strategies to a corresponding sequence of clock prices $p\left(A^t\right)$ that implements α. Importantly, our construction has the properties for each i and any legal history A^t of the auction that (a) $\max \hat{B}_i\left(A^t\right) = p_i\left(A^{t-1}\right)$ for every $i \in A_t$, $\min \hat{B}_i\left(A^t\right) > p_i\left(A^{t-1}\right)$ for every $i \in N \backslash A_t$, and $\alpha\left(b\right) \subseteq A_t$ for all $b \in \hat{B}\left(A^t\right)$, and (b) the strategy sets $\hat{B}\left(A^t\right)$ are determined by iterated deletion of dominated and equivalent strategies in a particular order for *any* generic value profile v such that $[v_i^+ \leq p_i\left(A^{t-1}\right)$ if and only if $i \in A_t]$. We establish properties (a) and (b) by induction.

We initialize the construction with clock prices $p\left(N\right) = \max B$ and sets of profiles $\hat{B}\left(N\right) = B$. For each clock round $t = 1, 2, 3, ...$, given any legal history A^t and previously determined strategy profiles $\hat{B}\left(A^t\right)$, we build $p\left(A^t\right)$ and $\hat{B}\left(A^{t+1}\right)$ as follows. Within each "clock" iteration t, we nest a second iteration employing a dummy variable $\bar{B}$. Initialize $\bar{B} = \hat{B}\left(A^t\right)$. Check whether there is some $i \in A_t$ and $b_i, b_i' \in \bar{B}_i$ such that $b_i' > b_i$ and $\alpha\left(b_i', b_{-i}\right) = \alpha\left(b_i, b_{-i}\right)$ for all $b_{-i} \in \bar{B}_{-i}$. If there is, we delete b_i from $\bar{B}_i$. Notice that the bid b_i is dominated by or equivalent to b_i' for all value profiles v (it wins against the same profiles b_{-i} and earns a higher price when it wins), so this step deletes only equivalent or weakly dominated strategies. Repeat this step to further trim $\bar{B}$ until the checking step indicates that no such qualifying bids $b_i' > b_i$ remain.

We claim that, after maximal trimming of $\bar{B}$, either all remaining strategy profiles lead to the same winners (i.e. $\alpha\left(\bar{B}\right) = \{A_t\}$) or else there exists an agent $i \in A_t$ for whom the bid $p_i\left(A^{t-1}\right) = \max \bar{B}_i$ always loses, that is, $i \in N\backslash\alpha\left(p_i\left(A^{t-1}\right), b_{-i}\right)$ for all $b_{-i} \in \bar{B}_{-i}$. To establish this claim, we use the inductive property and the assumption of dominance solvability for the game with a value profile v satisfying $v_i^+ < \min \bar{B}_i$ for $i \in A_t$ (so that agents $i \in A_t$ always strictly prefer to win) and $v_i^+ > p_i\left(A^{t-1}\right)$ for $i \in N\backslash A_t$ (the remaining agents prefer to lose). First, by inductive property (b), iterated deletion of dominated and equivalent strategies for v yields the sets $\hat{B}\left(A^t\right)$. Next, given such sets, if there are any two bids $b_i, b_i' \in \hat{B}_i\left(A^{t-1}\right)$ such that $b_i' > b_i$ and b_i is dominated by b_i', then (by monotonicity) both win against the same set of opposing bid profiles b_{-i} and hence (by non-bossiness) lead to the same allocations. Bids b_i that are dominated in this way are eliminated by the iterative "pruning" described above, at the end of which none such remain in $\bar{B}$. Hence, unless there is a unique set of winners ($\alpha\left(\bar{B}\right) = \{A_t\}$), dominance solvability for value profile v implies that there is another dominance relation to be found: there exists at least one active bidder $i \in A_t$ and bids $b_i, b_i' \in \bar{B}_i$ with with $b_i' < b_i$ such that b_i is dominated by b_i'. Given v, such dominance is possible only if b_i never wins, which by monotonicity implies our claim that $p_i\left(A^{t-1}\right) = \max \bar{B}_i \geq b_i$ never wins.

For iteration t of the clock auction, we reduce the price to the identified bidder i by letting $p_i\left(A^t\right) = \max\left(\bar{B}_i\backslash\left\{p_i\left(A^{t-1}\right)\right\}\right)$ and $p_j\left(A^t\right) = p_j\left(A^{t-1}\right)$ for every bidder $j \in N\backslash\{i\}$. The clock auction and the strategy sets for the next round are then updated as follows. If bidder i accepts the reduced clock price $p_i\left(A^t\right)$ at iteration t, we let $A_{t+1} = A_t$ and $\hat{B}_i\left(\left(A^t, A_t\right)\right) = \bar{B}_i\backslash\left\{p_i\left(A^{t-1}\right)\right\}$. If, instead, bidder i quits, we let $A_{t+1} = A_t\backslash\{i\}$ and $\hat{B}_i\left(\left(A^t, A_t\backslash\{i\}\right)\right) = \left\{p_i\left(A^{t-1}\right)\right\}$. For all bidders $j \in N\backslash\{i\}$, regardless of i's decision to accept or reject, we let $p_j\left(A^t\right) = p_j\left(A^{t-1}\right)$ and $\hat{B}_j\left(\left(A^t, A_t\right)\right) = \hat{B}_j\left(\left(A^t, A_t\backslash\{i\}\right)\right) = \bar{B}_j$. This guarantees that property (b) extends to both history $\left(A^t, A_t\right)$ and $\left(A^t, A_t\backslash\{i\}\right)$.

To see that with this construction, property (a) also extends from t to $t+1$, observe that it suffices to check the property for the bidder i whose price is changed. If $v_i^+ \leq p_i\left(A^{t-1}\right)$, then the bidder remains active and $\max \hat{B}_i\left(A^{t+1}\right) = p_i\left(A^t\right)$, as specified by the inductive property. Otherwise, $v_i^+ > p_i\left(A^{t-1}\right)$, the bidder quits and $\hat{B}_i\left(A^{t+1}\right) = \left\{p_i\left(A^t\right)\right\}$, so $\min \hat{B}_i\left(A^{t+1}\right) > p_i\left(A^t\right)$. By (a), the clock auction with cutoffs b leads to the outcome $\alpha(b)$.

■

Here are two examples of non-bossy allocation rules:

Example 20 (Optimization) *Letting* $F \subseteq 2^N$ *be the feasible set as in Example 4, the optimizing allocation rule is given by*

$$\alpha(b) \in \arg\min_{A \in F} \sum_{i \in A} b_i.$$

It is easy to see that, if B *rules out ties (so* $\arg\min$ *is always single-valued), optimizing allocation rules are non-bossy, because:*

$$\text{For } i \notin \alpha(b_i, b_i) \cup \alpha(b_i', b_i) \text{ we have } \alpha(b_i, b_i) = \arg\min_{A \in F : i \notin A} \sum_{j \in A} b_j = \alpha(b_i', b_i)$$

$$\text{For } i \in \alpha(b_i, b_i) \cap \alpha(b_i', b_i) \text{ we have } \alpha(b_i, b_i) = \arg\min_{A \in F : i \in A} \sum_{j \in A \setminus \{i\}} b_j = \alpha(b_i', b_i)$$

Example 21 (Fixed Scoring and Perfect Feasibility Checking) *Suppose we are in the setting of Example 4, that the feasible set* $F \subseteq 2^N$ *is comprehensive (as defined above), and that* $s_i^A(b_i, b_{N\setminus A}) = \begin{cases} \sigma_i(b_i) & \text{if } A \cup \{i\} \in F, \\ 0 & \text{otherwise,} \end{cases}$ *where the functions* $\sigma_i(b_i)$ *are increasing and positive-valued and there are no ties (so feasibility is always maintained). As observed above, every deferred acceptance procedure satisfies non-bossiness for the winning bids. To check that condition for rejected bids, too, suppose that given bid profile* b *agent* i*'s bid* b_i *is rejected in round* t *and agent* j*'s bid* b_j *is rejected in round* t *(hence* $A_t \setminus \{i, j\} \in F$*, and so by comprehensiveness* $A_t \setminus \{j\} \in F$*) but replace* b_i *with a bid* $b_i' < b_i$ *that is rejected in round* $t+1$*. In this case, bid* j *must be rejected in round* t *(so we must have*

$$\max_{k \in A_t \setminus \{i,j\} : A_t \setminus \{j,k\} \in F} \sigma_k(b_k) < \sigma_i(b_i') < \sigma_j(b_j).$$

After round $t+1$ *the heuristic is unaffected by the replacement. Iterating this argument, we see that any change in* b_i *that preserves this bid being losing will not affect the allocation produced by the heuristic.*

7 Comparisons to Properties of Other Auctions

The properties that we have derived for deferred acceptance auctions are not shared by other classes of auctions that have received close attention. Below are some examples to show that our findings do not apply to auctions in which winners are selected using either optimization or a greedy-acceptance heuristic.

7.1 Auctions Using Optimization

An optimizing allocation rule minimizes the total social cost subject to a feasibility constraint. Letting $F \subseteq 2^N$ be the feasible set as in Example 4, an optimizing rule solves

$$\alpha(b) \in \arg\min_{A \in F} \sum_{i \in A} b_i.$$

This is a monotonic allocation rule, and if $B_i = (0, +\infty)$ then the threshold prices are Vickrey prices - the agent is paid the externality his inclusion creates on the other agents:

$$p_i(b_{-i}) = \min_{A \in F : A \subseteq N \setminus \{i\}} \sum_{j \in A} b_j - \sum_{j \in \alpha(b) \setminus \{i\}} b_j$$

(so that his surplus $p_i(b_{-i}) - b_i$ captures the entire social cost savings due to his participation).

In some circumstances, the optimizing allocation rule and Vickrey prices can be computed with a deferred-acceptance heuristic or clock auction (ignoring any computational challenges that this might involve). This is determined by properties of the feasible set F. For example, when F is a comprehensive set and $\min B_i > 0$ for each i, the range of α has "no disposal" (as defined above). If α also satisfies substitutes, then by Proposition 14 allocation rule α is implementable by a clock auction or a deferred-acceptance heuristic when bid spaces B_i are finite. (See Bikhchandani et al. (2011) for conditions on F for an optimizing allocation rule to satisfy substitutes; see also Ausubel (2004) and de Vries and Vohra (2007) for earlier examples of settings in which optimizing allocation rules can be implemented via clock auctions.) In this case, paid-as-bid equivalence also holds. Bernheim and Whinston (1986) had shown payoff equivalence between Vickrey and paid-as-bid auctions when bidders are substitutes using a coalition-proofness refinement to select among Nash equilibrium. Our analysis finds the same conclusion under different assumptions and conditions. We use either iterated dominance or undominated Nash equilibrium to select a Nash equilibrium and we allow a wide range of heuristic allocation rules with substitution, but we limit attention to environments with single-minded bidders.

For an example in which the substitutes condition does not hold and so an optimizing α cannot be implemented via a deferred-acceptance heuristic, consider the following:

Example 22 $N = \{1, 2, 3\}$ *and* $F = \{\{1, 2\}, \{3\}\}$. *Intuitively, the structure of* F *makes bidders* 1 *and* 2 *complementary. In this case,* $\alpha(b) = \{1, 2\}$ *if* $b_1 + b_2 < b_3$ *and* $\alpha(b) = \{3\}$ *if* $b_1 + b_2 > b_3$. *In any deferred-acceptance heuristic, the first bid to be rejected can be based only on pairwise comparisons of bids, so it cannot be generally consistent with the preceding inequalities.*

Observe, too, that the Vickrey auction implementing this allocation rule does not satisfy either weak group strategy-proofness or paid-as-bid equivalence. For example, when $b_1 + b_2 < b_3$, the Vickrey prices are $p_1(b_2, b_3) = b_3 - b_2$ and $p_2(b_2, b_3) = b_3 - b_1$. Then the two winners have a strictly improving coalitional deviation in which they bid $b_1 < v_1$, $b_2 < v_2$ such that $b_1 + b_2 < v_3$: both still win but each is paid strictly more. Also, note that the sum of Vickrey prices is $2b_3 - b_1 - b_2 > b_3$, but in the corresponding paid-as-bid auction, there cannot be a Nash equilibrium in which bidders 1 and 2 win and are paid a total of $b_1 + b_2 > v_3$, since then bidder 3 would deviate to undercut them. (In fact, in all the Nash equilibria in which bidder 3 uses undominated strategies and bidders 1 and 2 win, they together pay $b_1 + b_2 = v_3$. These outcomes have been identified by Bernheim and Whinston (1986).) So, the Vickrey mechanism appears "too expensive" in this case relative to optimization with paid-as-bid pricing. One solution that has been proposed to the problem of Vickrey auctions' excessive costs (insufficient revenues) is "core-selecting auctions" (Day and Milgrom 2008), which sacrifice strategy-proofness even for single-minded bidders. Deferred-acceptance heuristics offer a possible alternative way to reduce costs (increase revenues), which preserves strategy-proofness.

7.2 Auctions Using Greedy-Acceptance Heuristics

To compare to the greedy-acceptance heuristics auctions of LOS, consider again Example 22. For illustration, let a bidder's score be its bid, so the heuristic iterates accepting the highest bid that is still feasible. If we break ties in favor of lower-numbered agents, we have $\alpha(b) = \{b_1, b_2\}$ if $\min\{b_1, b_2\} \geq b_3$, and $\alpha(b) = \{b_3\}$ otherwise. Suppose bid spaces are $b_i = [0, \bar{b}]$. The threshold payments for the reverse auction are as follows: First, if $\alpha(b) = \{b_1, b_2\}$, then for $i = 1, 2$, $p_i = \bar{b}$ if $b_{-i} \leq b_3$ and $p_i = b_3$ otherwise. Second, if $\alpha(b) = \{b_3\}$, then $p_3 = \min\{b_1, b_2\}$.

Observe that this allocation rule cannot be implemented with a deferred-acceptance heuristic or a descending clock auction, since the allocation is completely determined by the single "best" (lowest) bid while the first step of the deferred acceptance heuristic is determined by the single worst bid according to some criterion. One might conjecture that greedy acceptance heuristics could instead be matched with an ascending clock auction, but that fails, too, because when bidder 3 exits first, the allocation is determined to be $\alpha(b) = \{b_1, b_2\}$ but the prices to the winners are not yet determined.[15]

Next, observe that the greedy-acceptance threshold auction fails weak group strategy-proofness. For example, if $v_1, v_2 > v_3 > 0$, bidders 1 and 2 could jointly deviate to bid $b_1, b_2 < v_3$, which will give each of them threshold prices of $\bar{b}$.

Finally, the threshold and paid-as-bid auctions based on the greedy-acceptance heuristic do not have the outcome equivalence properties described above. For suppose that $v_1, v_2 < v_3$. In the threshold auction, bidders 1 and 2 win and their threshold payments are both $\bar{b}$, but a paid-as-bid auction with complete information cannot have a pure Nash equilibrium in which bidders 1 and 2 win and both get paid above v_3, since then bidder 3 would deviate to undercut them both and win. So a greedy-acceptance heuristic with threshold payments is more expensive in this case than any pure Nash equilibrium of its paid-as-bid counterpart.

8 Conclusion

8.1 Clock Auctions vs. Sealed-Bid Heuristics

We have identified several attractive properties that are satisfied by deferred-acceptance threshold auctions, whether implemented as sealed-bid auctions or using an equivalent clock auctions. In richer practical settings, however,

[15] A traditional purpose of a clock auction is to economize on information transmission or conceal some information, and accordingly we require that that clock-auction prices stop changing once the allocation has been determined. Without this condition, any allocation rule can be implemented with a clock auction, simply by running all the prices down to elicit complete information about cutoffs from all bidders and determining the allocation as a function of those.

the two kinds of auctions are not be equivalent. Clock auctions may often be preferred, for the following reasons:

1. Bidders for whom it is costly to figure out or to disclose their exact values will find it easier to participate in a clock auction than in an equivalent sealed-bid auction. Indeed, in a clock auction, losers do not need to figure out or disclose their exact values - they only need to establish that these values are below the final price.

2. For bidders who do not understand the auction or do not trust the auctioneer to follow the rules, a clock auction is preferable: the auctioneer only needs to promise that prices can only go down and never up, which makes the optimality of truthful bidding obvious for private-value bidders. In contrast, in equivalent sealed-bid auction, bidders may not understand the argument that truthful bidding is optimal, or may not trust the auctioneer not condition the outcome on the revealed in a way that makes truthful bidding suboptimal. Indeed in experiments, bidders were found to shade their bids in simple sealed-bid threshold-price auctions while bidding truthfully in strategically equivalent clock auctions (Kagel et al. 1987). These concerns about transparency and commitment are likely to be much more prominent for more complicated auctions.

3. When bidders' values are not private but interdependent, cutoff strategies will no longer be best responses for them in clock auctions some information disclosure: instead, they will condition their exit decisions on the disclosed information. It may be possible to tune decisions about information disclosure to enhance efficiency or other objectives of the auction.

Against these advantages, a possible disadvantage of a clock auction with small price decrements could take many rounds to conclude, raising costs to both the auctioneer and the bidders.

8.2 Algorithmic Properties

Despite our suggestion that the deferred acceptance auctions might be useful for complex resource allocation, this paper has focused on the game-theoretic properties of the auctions while mostly ignoring the performance of the algorithms themselves. The exception is our theorem asserting that deferred acceptance heuristics can implement allocation rules with the substitutes property.

We have not examined issues such as computation time, nor have we studied the approximation error in comparison to full optimization. There may well be a trade-off here, as particular scoring functions or price computations may be more easily computed than others, and may lead to different levels of efficiency or revenue. There has been some other literature (e.g. Mehta et al. 2009) identifying cases in which simple deferred-acceptance heuristics and clock auctions can achieve approximate efficiency.

8.3 Multidimensional Bidders

We have restricted attention to the case of single-minded bidders, while in practice bidders are often interested in selling or buying various goods or bundles and have private values for those packages. For example, in the FCC problem that has motivated us, some bidders, as an alternative to selling their broadcast stations to go off-air, may also consider switching to a lower, less-congested band. Also, some broadcasters own multiple stations and may wish to contemplate which subset of stations to sell.

Clock auctions have been used and studied for such cases, but truthful bidding is not necessarily optimal for multi-minded bidders. Milgrom (2000) and Gul and Stachetti (2000) examine simple heuristic clock auctions (e.g. forward auctions, raise prices for overdemanded goods). If bidders bid "straightforwardly," then such auctions achieve efficiency under "substitute" conditions; however, they are not strategy-proof for bidders who are not single-minded.

Some clock auctions do provide incentives for truthful bidding for multi-minded bidders: Auctions that are based on Vickrey: Ausubel, Bikhchandani – de Vries – Vohra can be done under "substitution" conditions. However, they may require large computational burden.

A third possibility is to use auctions that are both strategy-proof for multi-minded bidders and simple to compute, but do not guarantee efficiency. E.g., Bartal et al. (2003) propose a clock auction in which each bidder is only asked once, with prices to a bidder depending on what the previous bidders chose. These auctions do not achieve efficiency, but some approximation thereof.

References

[1] Abdulkadiroglu, A., P. Pathak, A. Roth and T. Sonmez (2005), "The Boston Public School Match," *American Economic Review*, 95 (2), 368-371.

[2] Ausubel, L.M. (2004) "An efficient ascending-bid auction for multiple objects," *American Economic Review*, 94 (5), 1452–1475

[3] Babaioff, M., and L. Blumrosen (2008), "Computationally-feasible truthful auctions for convex bundles," *Games and Economic Behavior* 63(2), 588–620

[4] Bartal, Y., R. Gonen, and N. Nisan (2003), "Incentive compatible multi unit combinatorial auctions," *Proceedings of the 9th conference on Theoretical aspects of rationality and knowledge*, Pages 72 - 87

[5] Bernheim, D., and M. Whinston (1986), "Menu Auctions, Resource Allocation, and Economic Influence," *Quarterly Journal of Economics* 101(1): 1-31.

[6] Bikhchandani, S., S. de Vries, S., J. Schummer, and R.V. Vohra (2011), "An Ascending Vickrey Auction for Selling Bases of a Matroid," *Operations Research* 59 (2), 2011, pp. 400–413

[7] Day, R. W., and P. Milgrom (2008), "Core-Selecting Package Auctions." *International Journal of Game Theory*, 36, 2008, 393–40

[8] de Vries, S., J. Schummer, and R.V. Vohra (2007), "On ascending Vickrey auctions for heterogeneous objects," *Journal of Economic Theory* 132(1), 95–118

[9] Ensthaler, L., and T. Giebe (2009), "Subsidies, Knapsack Auctions and Dantzig's Greedy Heuristic," SFB/TR 15 Discussion Paper No. 254, DIW Berlin Discussion Paper No. 880

[10] Ensthaler, L., and T. Giebe (2010), "A dynamic auction for multi-object procurement under a hard budget constraint," SSRN Working Paper

[11] Fujishima, Y., D. McAdams, and Y. Shoham (1999), "Speeding up ascending-bid auctions," Proceedings of the International Journal Conference in Artificial Intelligence, 554-559

[12] Gul, F., and E. Sachetti (2000), "The English auction with differentiated commodities," *Journal of Economic Theory* 92(1), 66–95

[13] Hayek, F.A. (1945): "The Use of Knowledge in Society," *American Economic Review* 35, 519-30.

[14] Hatfield, J., and P. Milgrom (2005), "Matching with Contracts," *American Economic Review*

[15] Hurwicz, L. (1977): "On the Dimensional Requirements of Informationally Decentralized Pareto-Satisfactory Processes," in K.J. Arrow and L. Hurwicz, eds., *Studies in Resource Allocation Processes*, 413-424, New York: Cambridge University Press.

[16] Juarez, R. (2009), "Prior-free cost sharing design: group strategyproofness and the worst absolute loss," in *Social Computing and Behavioral Modeling,* M. Young, J. Salerno and H. Liu (eds), Springer.

[17] Kagel, J. H., R.M. Harstad, and D. Levin (1987), "Information impact and allocation rules in auctions with affiliated private values: a laboratory study," *Econometrica* 55, 1275-1304.

[18] Kelso, A., and V. Crawford (1982), "Job Matching, Coalition Formation, and Gross Substitutes," *Econometrica* 50(6), 1483-1504.

[19] Lehmann, D., L.I. O'Callaghan, Y. Shoham (2002), "Truth Revelation in Approximately Efficient Combinatorial Auctions," *Journal of the ACM* 49(5), 577-602.

[20] Marx, L.M., and J.M. Swinkels (1997), "Order Independence for Iterated Weak Dominance," *Games and Economic Behavior* 18(2), 219–245.

[21] Mehta, A., T. Roughgarden, and M. Sundararajan (2009), "Beyond Moulin Mechanisms," *Games and Economic Behavior* 67(1), 125–155.

[22] Milgrom, P. (2000), "Putting Auction Theory to Work: The Simulteneous Ascending Auction," Journal of Political Economy 108(2), pp. 245-272.

[23] Mount, K., and S. Reiter (1974): "The Information Size of Message Spaces," *Journal of Economic Theory* 28, 1-18.

[24] Mu'alem, A., and N. Nisan (2008), "Truthful approximation mechanisms for restricted combinatorial auctions," *Games and Economic Behavior* 64(2), 612–631

[25] Moulin, H. (1999), "Incremental cost sharing: Characterization by coalition strategy-proofness," *Social Choice and Welfare* 16, 279-320.

[26] Nisan, N., and I. Segal (2006), "The communication requirements of efficient allocations and supporting prices," *Journal of Economic Theory* 129 (1), 192-224

[27] Papadimitriou, C., and Y. Singer (2010), "Budget Feasible Mechanisms," arXiv:1002.2334 preprint

[28] Segal, I. (2007), "The Communication Requirements of Social Choice Rules and Supporting Budget Sets," *Journal of Economic Theory* 136, pp. 341-378