Deferred-Acceptance Auctions and Radio Spectrum Reallocation*

Paul Milgrom[†] Ilya Segal[‡]

September 11, 2014

Abstract

Deferred-acceptance auctions choose allocations by an iterative process of rejecting the least attractive bid. These auctions have distinctive computational and incentive properties that make them suitable for application in some challenging environments, such as the planned US auction to repurchase television broadcast rights. For any set of values, any deferred acceptance auction with "threshold pricing" is weakly group strategy-proof, can be implemented using a clock auction, and leads to the same outcome as the complete-information Nash equilibrium of the corresponding paid-as-bid auction. A paid-as-bid auction with a non-bossy bid-selection rule is dominance solvable if and only if it is a deferred acceptance auction.

*The authors thank Alexey Kushnir as well as Mohammad Akbarpour, Piotr Dworczak, Tim Roughgarden, Inbal Talgam Cohen, Joshua Thurston-Milgrom, and Alex Wolitzky for their comments and suggestions.

[†]Department of Economics, Stanford University, Stanford, CA 94305. The authors are members of the Auctionomics' team advising the US Federal Communications Commission concerning the design of the US "incentive auction." This research analyzes some properties of proposals to which the authors have contributed.

[‡]Segal gratefully acknowledges the support of the Toulouse Network for Information Technology.

1 Introduction

We study a class of heuristic auctions for computationally challenging resource allocation problems. The application that motivates this study arises from the US government's effort to reallocate channels currently allocated for television broadcasting to use for wireless broadband services. This reallocation involves purchasing television broadcast rights from some TV stations, reassigning ("repacking") the remaining over-the-air broadcasters into a smaller set of channels, using the cleared spectrum to create licenses suitable for use in wireless broadband, and selling those licenses to cover the costs of acquiring broadcast rights.[1] The repacking of about 2,000 TV stations must be done in a way that satisfies about 130,000 interference constraints, which dictate which pairs of stations could not be assigned to the same or adjacent channels. Even determining whether a given set of broadcasters can be feasibly assigned to the available set of channels is a computationally challenging problem, equivalent to the NP-hard "graph coloring problem" (see Aardal et al. (2007)). The problem of selecting the set of stations to maximize the total broadcast value in the remaining TV spectrum is an even harder computational problem, which cannot be solved exactly in reasonable time using today's state-of the art algorithms and hardware. This creates a challenge for designing a strategy-proof auction: for example, the approximation error in computing Vickrey prices would be quite large relative to the bids of individual bidders, potentially creating large incentives for misrepresenting preferences.[2]

[1] Middle Class Tax Relief and Job Creation Act of 2012, Pub. L. No. 112-96, §§ 6402, 6403, 125 Stat. 156 (2012). The legislation aims to make it possible to reallocate spectrum from a lower-valued use to a higher-valued one, generating billions in net revenues for the federal government.

[2] The problem of designing computationally feasible economic mechanisms is studied in the field of "Algorithmic Mechanism Design" (Nisan and Ronen (1999)). While economists have long been concerned about the computational properties of economic allocation mechanisms (e.g., see Hayek (1945)), the formal economic literature has focused on modeling communication costs (e.g., Hurwicz 1977, Mount and Reiter 1974, Segal 2007), which

Our proposed solution is to use the class of "deferred-acceptance" (DA) auctions, which generalize the idea of the eponymous matching algorithms of Gale and Shapley (1962) and Kelso and Crawford (1982) and related clock auction designs, such as the simultaneous multiple-round auctions used for radio spectrum sales. These auctions are iterative processes in which the highest scoring (and hence least attractive) bids are rejected in each round and the bids left standing at the end are accepted.

For concreteness, we focus on the procurement setting (the modification to the selling setting is straightforward). We also restrict attention to settings in which each bidder is "single-minded," meaning that it offers one good (or bundle of goods) for sale and cares only about winning or losing and the amount of its compensation. A DA auction in this setting is based on the following kind of DA algorithm: All the bids are initially active. In each iteration, each bid is assigned a non-negative "score" according to a scoring function whose arguments include the bid amount, the bidder's identity, the remaining set of active bidders, and the amounts and identities of the previously rejected bids. If no active bid has a positive score, the algorithm terminates and the remaining active bids become winning. Otherwise, the algorithm rejects the active bid with the highest non-zero score, removing it from the active set, and iterates. Different scoring rules lead to different DA algorithms.

Deferred-acceptance auctions are sealed-bid mechanisms that select winning bids using a DA algorithm and make payments only to winning bidders. In particular, we consider auctions that pay each winning bidder his *threshold price*, defined as the highest bid he could have made while still winning, holding all the other bids fixed. It is routine to show that any *deferred-acceptance threshold auction*, which uses a DA algorithm to select winners and pays each winner its threshold price, is a strategy-proof auction.

By selecting a suitable scoring rule, the DA auctions can achieve a variety

---

are trivial in the setting of single-minded bidders considered in this paper.

of different goals and can be shown to nest several previously studied auctions. If the goal is to implement an efficient allocation subject to constraints that have a matroid structure (informally, when adding one bidder displaces at most one other bidder), that can be achieved exactly using the "standard" DA algorithm of Bikhchandani et al. (2011), in which all active bids are scored by their bid values. In settings with more complex constraints, DA algorithms may not be able to achieve exact efficiency, but may still approximate it fairly closely. For example, a DA auction can implement Dantzig's (1957) greedy algorithm for the "knapsack problem," in which there is a constraint on the total "volume" of the rejected bids. For the FCC's repacking problem, our simulations using appropriate scoring for stations based on their interference constraints have generated outcomes, for sample station values, that are within percentage points of exact efficiency.[3]

In some cases, efficiency is not the only or even the primary requirement. For example, the auctioneer may face a hard budget constraint (as the FCC may in its repacking problem). DA threshold auctions can respect such constraints, nesting the budget-constrained mechanisms previously studied by McAfee (1992), Moulin (1999), Juarez (2007), Mehta et al. (2007), Singer (2010), and Ensthaler and Giebe (2014). Another potentially important goal may be to minimize the procurement cost. If bidder values are drawn independently from known distributions, then bids in DA auctions can be scored according to the corresponding virtual values of Myerson (1991). In that case, if the constraints have a matroid structure, the upshot is an expected-cost minimizing auction. Alternatively, if bidder values are drawn i.i.d. from a single *unknown* distribution, then expected costs can be reduced by scoring the still-active bids using already-rejected bids (similarly to the "yardstick

[3]The auction requires checking in every round whether any given active station could be feasibly rejected and added to the already-assigned stations with possible repacking. While this is an NP-hard computational problem, it turns out that this computation could be sped up to a matter of seconds. When the computation does time out, which happens occasionally, we treat the answer as infeasibility and so cannot reject the bidder.

competition" in Segal (2003)).

The main theoretical contribution of the paper lies in establishing several nice properties shared by all DA threshold auctions. First, we establish an equivalence between the sets of DA threshold auctions and descending "clock" auctions. Intuitively, similarly to a DA auction, a clock auction iteratively rejects the bids that are currently the least attractive to the auctioneer, but unlike the DA auction it permits the rejected bidders to remain active by improving their bids by some minimal bid increment. Formally, we show that for every sealed-bid DA threshold auction with finite bid spaces, there exists an equivalent descending clock auction that generates the same allocation and payments, and vice versa. Given the theoretical equivalence, clock auctions may be preferable for many practical applications, because they (i) make strategy-proofness self-evident even for bidders who misunderstand or mistrust the auctioneer's calculations (see Kagel (1987) for experimental evidence), (ii) do not require winning bidders to reveal or even figure out their exact values, and (iii) permit information feedback so as to better aggregate common-value information.

A second nice property is that DA threshold auctions and the corresponding clock auctions are weakly group strategy-proof: no coalition of bidders has a joint deviation from truthful bidding that is strictly profitable for all members of the coalition. For the clock auctions, this property holds regardless of what information is disclosed to bidders during the auction. This result extends weak group strategy-proofness results obtained earlier by Moulin (1999), Mehta et al. (2007), and Hatfield and Kojima (2009).

A third set of results stems from our investigation of the acquisition cost achieved by DA threshold auctions compared to similar paid-as-bid mechanisms. The challenges in making this comparison are two. First, our setting is too complicated to allow Bayesian equilibrium analysis of paid-as-bid auctions, so we are limited to full-information equilibrium. Second, even with full-information equilibrium, paid-as-bid auctions based on optimization are known to have many different Nash equilibrium outcomes, so we need to explore that possible multiplicity for DA auctions as well. Our next two findings concern equilibrium properties of DA paid-as-bid auctions that are surprisingly similar to those of paid-as-bid auctions for a single item. These familiar properties have not been reported for other multi-item auctions, and our third finding explains why: *every multi-item paid-as-bid auction with these familiar properties is a DA-auction.* The findings are as follows. First, every paid-as-bid DA auction has a full-information Nash equilibrium that generates the same outcome as the threshold auction using the same DA algorithm. Second, when we limit attention to "non-bossy" DA algorithms [4] and to generic bidder values, the DA paid-as-bid auction has a *unique* undominated Nash equilibrium outcome and a unique outcome that survives iterated elimination of weakly dominated strategies, and both coincide with the dominant-strategy outcome of the corresponding DA threshold auction. Third, these properties are enjoyed *only* by DA auctions: any non-bossy paid-as-bid auction that is generically dominance solvable is a DA auction.

To highlight further how special the properties of DA auctions are, we compare the properties of other auctions with single-minded bidders in which winning bids are selected in two other familiar ways. First, some auctions are based on optimization, in which the collection of winning bids is the one that minimizes the total bid. When the constraints have a matroid structure that ensures that bidders are substitutes, these auctions can be implemented as DA auctions and therefore have all the nice properties described above, but otherwise they do not: they (i) cannot be implemented as clock auctions,[5] (ii) are not weakly group strategy-proof, and (iii) yield threshold prices (Vickrey prices) that result in a higher procurement cost than

[4] These are algorithms for which a bidder cannot alter any part of the allocation without changing his own allocation.

[5] Ausubel and Milgrom (2002) propose a "cumulative-offer" clock auction which implements efficiency, but this auction requires the ability to "recall" exited bids, and this violates strategy-proofness.

the full-information Nash equilibria of the corresponding paid-as-bid auction (see Ausubel and Milgrom (2006), Bernheim and Whinston (1986)).[6] Second, Lehmann et al. (2002) introduced a class of auctions that select winners using a "greedy-acceptance" algorithm, meaning that which bids are accepted iteratively until certain constraints are satisfied.[7] In general, these greedy-acceptance auctions permit the calculation of threshold prices that ensure strategy-proofness, and they can achieve good approximations of efficiency in some interesting, computationally challenging applications. Furthermore, when the constraints have a matroid structure (meaning that bidders are substitutes for the auctioneer), there exists a greedy-acceptance algorithm that coincides both with the Vickrey auction and with its corresponding DA auction. However, as we show in Appendix A by means of an example, greedy-acceptance algorithms do not *generally* enjoy any of the other nice properties described above: they (i) cannot be implemented as clock auctions, (ii) are not weakly group strategy-proof, and (iii) can result in a higher procurement cost than some full-information undominated Nash equilibrium of the corresponding paid-as-bid auction.

The paper is organized as follows: Section 2 illustrates a simple deferred-acceptance auction in a three-bidder example and contrasts its properties with those of the Vickrey auction. Section 3 introduces our general model and characterizes strategy-proof sealed-bid auctions in our setting. Section 4 introduces general deferred-acceptance auctions. Section 5 provides several examples of useful DA auctions. Section 6 introduces descending-clock auctions in which winners are paid their final clock prices and shows the equiv-

[6] Indeed, this high costs/low revenue problem of the Vickrey auction is the motivation for the proposed "core-selecting auctions" (Day and Milgrom 2008), which sacrifice strategy-proofness in order to increase revenue. The present paper proposes a different solution to the problem, preserving strategy-proofness but sacrificing the efficiency of outcome according to the bids. (Of course, core-selecting auctions only guarantee that the outcome is efficient according to the bids, and so efficiency of outcome according to true values is not ensured given the gaming opportunities generated by these auctions.)

[7] For other examples of such auctions, see also Mu'alem and Nisan (2008), Babaioff and Blumrosen (2008), and the references therein.

alence between these auctions and DA threshold auctions. Section 7 shows that descending clock auctions (and the equivalent threshold-price DA auctions) are weakly group strategy-proof. Section 8 analyzes paid-as-bid DA auctions, deriving their properties described above. Section 9 concludes with a discussion of how some of these results might be used for the FCC's auction design problem.

2 A Simple Example

In this section, we compare a DA auction with a Vickrey auction in a simple setting to illustrate some advantages of DA auctions beyond computational simplicity. In our example, the bidders are three TV stations (labelled 1,2,3) bidding to relinquish broadcast rights in a reverse auction, with only a single channel available to assign the losing bidders. Stations 1 and 2 are "peripheral" stations that can be assigned to broadcast on the same channel, but either of them would interfere with the "central" station 3. Thus, the auctioneer needs to choose between clearing the desired spectrum by acquiring station 3 (assigning both stations 1 and 2 to the remaining channel) and doing it by acquiring stations 1 and 2 (assigning station 3).[8] Suppose the three stations bid respective amounts b_1, b_2, b_3 to be paid for relinquishing their broadcast rights.

In this setting, the Vickrey auction implements the efficient allocation, which is to acquire stations 1 and 2 when $b_1 + b_2 \leq b_3$ and station 3 otherwise. (For concreteness, we break ties in favor of the lower-numbered bidders.) If the auction acquires stations 1 and 2, the Vickrey payments to them are $b_3 - b_2$ and $b_3 - b_1$, respectively. A first concern about the Vickrey auction is that it allows stations 1 and 2 to raise each other's payments by reducing their bids, and in particular to each obtain the maximal payment of b_3 by

[8] It is also feasible, but clearly suboptimal, to acquire and clear station combinations {1,3}, {2,3}, and {1,2,3}.

both bidding zero. This demonstrates the fact that the Vickrey auction is not (even weakly) group strategyproof. Second, the total payment to bidders 1 and 2, which equals $2b_3 - b_1 - b_2$, exceeds station 3's value b_3. In contrast, in the *paid-as-bid* auction in which the auctioneer selects the cost-minimizing combination of bids (a "menu auction" of Bernheim and Whinston (1986)), there is no full-information Nash equilibrium in which the total payment to bidders 1 and 2 exceeds bidder 3's value, since bidder 3 could have profitably deviated by undercutting $b_1 + b_2$. In the general setting, Vickrey payments are sometimes higher, but never lower, than those in the full-information equilibrium of the corresponding paid-as-bid ("menu") auction (Ausubel and Milgrom (2006)).

The DA threshold auctions proposed in this paper are weakly group strategyproof, and do not cost more than the corresponding paid-as-bid equilibrium cost, but sacrifice the requirement of always choosing the efficient allocation given the bids. (In contrast, the "core-selecting auctions" (Day and Milgrom (2008)) address the revenue shortfall problem by modifying transfers but insisting on the efficient allocation, which sacrifices strategyproofness.)

A DA threshold auction in this setting begins with all bidders active and, in each round, irreversibly rejects the bidder with the highest positive score, where any bid that cannot be feasibly rejected gets a score of zero. For a simple example, suppose that bidders 1 and 2 are given scores equal to their bids b_1, b_2 so long as they are feasible, while bidder 3 is given a score equal to b_3/w_3. (The factor $w_3 \geq 1$ is interpreted as the bidder's "volume", to account for the greater interference it creates.) With these scores, the DA algorithm acquires stations 1 and 2 if both $b_1, b_2 \leq b_3/w_3$ (so that station 3 is rejected in the first round), and acquires station 3 otherwise. To ensure strategyproofness, when stations 1 and 2 are acquired, each of them is paid b_3/w_3 (this is the station's "threshold price," which is its maximal bid that would have still been accepted given the bids of the others). When station 3 is acquired, it is paid its "threshold price" $w_3 \max\{b_1, b_2\}$.

In contrast to the Vickrey auction, the DA threshold auction is weakly group strategyproof, for if both bidders 1 and 2 both win, their bids do not affect each other's payments. The auction also satisfies a "revenue equivalence" property: in a paid-as bid auction with the same allocation rule, when the stations' true values satisfy $b_1, b_2 \leq b_3/w_3$, the unique full-information Nash equilibrium in weakly undominated strategies is for station 3 to lose by bidding its true value b_3, and for stations 1 and 2 to win by bidding their "threshold prices" b_3/w_3. Finally, the threshold-price DA auction has a simple descending-price clock auction implementation: we start a per-volume price p at a high level and let it descend continuously, permitting each bidder to exit at any point at which it is still feasible, and "freezing" the prices to those bidders that can no longer feasibly exit. (For example, if station 3 is the one that exits first, at the clock price $p = b_3/w_3$, then the other two stations are frozen at this price and the auction ends.) An advantage of this clock auction is that it makes the optimality of truthful bidding self-evident: a bidder could not benefit from either quitting at a price above his value or staying at a price below his value. In contrast, the Vickrey auction described above does not have a clock-auction implementation, bidders who are faced with its sealed-bid form may fail to understand the optimality of truthful bidding (see the experiments of Kagel et al. (1987)).

As described below, general DA auctions can use more general scoring functions than the one described above, which may also vary from round to round and may depend on the previously rejected bids. Nevertheless, all these auctions have the nice properties mentioned above. At the same time, none of these auctions would guarantee the achievement of an efficient allocation in this section's example. Intuitively, this is because any DA auction would reject the first bid based on the pairwise comparison of its score with those of the other bids, rather than on the comparison of $b_1 + b_2$ to b_3, which is required to guarantee efficiency.

In some important practical settings, guaranteeing efficiency may be less important than other objectives, for any of the following reasons: (1) Exact efficiency may be unachievable due to computational problems. (2) Efficiency may be undermined by bidder manipulation in mechanisms that are not strategy-proof (such as core-selecting auctions or approximate-Vickrey auctions) or not group strategy-proof (such as Vickrey auctions). (3) The auction designer may be more interested in other objectives, such as procurement cost, and may also face a budget constraint. In such settings, the designer may be well served by using one of the class of DA auctions studied in this paper.

3 Sealed-Bid Auctions

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 bundle of goods is "accepted") or "loses" (which means that his bid is rejected). We restrict attention to mechanisms, which we call "auctions," in which winning bidders receive payments but losing bidders do not.

The preferences of each bidder i depend on whether he wins or loses, and, if he wins, on the payment p_i. We assume that these preferences are strictly increasing in the payment and that there exists some payment v_i that makes him indifferent between winning and losing; we call v_i his "value".[9] The set of bidder i's possible values is $[0, \overline{v}_i]$ for each i.

We restrict each bidder to make bids in a finite set $B_i \subseteq \mathbb{R}_+$ and assume that $\max B_i > \overline{v}_i$. This ensures that a non-participation strategy is never better for bidder i than bidding $\max B_i$, so we can simplify our formulation by not allowing non-participation.

A sealed-bid auction takes a bid $b_i \in B_i$ from each bidder $i = 1, ..., N$. For every bid profile $b \in B = \Pi_i B_i$, the auction's *allocation rule* $\alpha : B \to 2^N$

[9] For unmixed outcomes, such a preference can be expressed by a quasilinear utility $p_i - v_i$ when the bidder wins and zero when he loses.

determines a set of winning bids $\alpha(b) \subseteq N$ and its *payment rule* $p : B \to \mathbb{R}^N$ specifies a corresponding payment profile, with $p_i(b) = 0$ for $i \in N \backslash \alpha(b)$ (so losing bidders are not paid). A *sealed-bid auction* is a triple $\langle B, \alpha, p \rangle$.

With finite bid sets, it is convenient to replace the usual notion of "truthful" bidding by a similar concept of "strategy-proofness" that applies even when truthful reporting is infeasible because $v_i \notin B_i$. According to our definition, an auction is strategy-proof if it is always optimal for a bidder to round up its value to the next allowable bid.

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 $v_i^+ \equiv \min\{b_i \in B_i : b_i > v_i\}$.*

The usual characterization of strategy-proof auctions applies in this setting. For completeness, we restate it here.

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 $i \in \alpha(b_i', b_{-i})$.*

Definition 3 *A sealed-bid auction $\langle B, \alpha, p \rangle$ is a* threshold auction *if and only if α is monotonic and the price paid to any winning bidder $i \in \alpha(b)$ satisfies the* threshold pricing *formula:*

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

With these definitions, a standard argument implies the following result.

Proposition 1 *A sealed-bid auction $\langle B, \alpha, p \rangle$ is strategy-proof if and only if it is a threshold auction.*[10]

[10]It is obvious that a threshold auction is strategy-proof. For the converse, apply the envelope theorem to see that the outcome function and strategy-proofness determine the bidders' payoffs up to a constant of integration determined. Since losers must be paid zero and have a payoff of zero, the strategy-proof payment function is unique.

4 Deferred-Acceptance Auctions

Given a set of bidders N and possible bid profiles B, a deferred-acceptance algorithm (DA) is specified by a collection of *scoring functions* $\left(s_i^A\right)_{A\subseteq N, i\in A}$, where for each $A \subseteq N$ and each $i \in A$, the function $s_i^A : B_i \times B_{N\setminus A} \to \mathbb{R}_+$ is nondecreasing in its first argument. The algorithm operates as follows. Let $A_t \subseteq N$ denote the set of "active bidders" at the beginning of iteration t. We initialize the algorithm with $A_1 = N$. In each iteration $t \geq 1$, 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 $\alpha(b) = A_t$; 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. In words, the algorithm iteratively rejects the least desirable (highest-scoring) bids until only zero scores remain. We say that the DA algorithm *computes* allocation rule α if for every bid profile $b \in B$, when the algorithm stops the set of active bidders is exactly $\alpha(b)$.

By inspection, every DA algorithm is finite and computes a monotonic allocation rule. Thus, we can define a *DA threshold auction* as a sealed-bid auction which computes its allocation using a DA algorithm and makes the corresponding threshold payments (1) to the winners. Like any threshold auction, each DA threshold auction is strategy-proof. Furthermore, the threshold prices can be computed in the course of the DA algorithm, by initializing the prices as $p_i^0 = \sup B_i$ for all i, and then updating them in each round $t \geq 1$ as follows:

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

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.

5 Examples of Deferred-Acceptance Auctions

Example 1 (Feasibility and Non-Wastefulness) *Suppose that it is only feasible for the auctioneer to acquire products from a subset of bidders* $A \in \mathcal{F}$*, where* $\mathcal{F} \subseteq 2^N$ *is a given family of sets, with* $N \in \mathcal{F}$ *(so that feasibility is achievable).*[11] *To ensure that the algorithm maintains feasibility, we require that* $s_i^A\left(b_i, b_{N\backslash A}\right) > 0$ only if $A \backslash \{i\} \in \mathcal{F}$*, and 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*,* $b_i, b_j, b_{N\backslash A}$*.*

Say that the family $\mathcal{F}$ *is* comprehensive *if* $A \in \mathcal{F}$ *and* $A \subseteq A'$ *imply* $A' \in \mathcal{F}$ *(i.e.,* $\mathcal{F}$ *has "free disposal" of winners). Say that the algorithm has* perfect feasibility checking *if* $s_i^A\left(b_i, b_{N\backslash A}\right) > 0$ if and only if $A \backslash \{i\} \in \mathcal{F}$ *– that is, it stops only when it is infeasible to reject any more active bids. If both conditions hold, then the algorithm has the desirable property of* non-wastefulness *– i.e., it stops at a* minimal *set* $A \in \mathcal{F}$*. (Indeed, if there is some* $A' \in \mathcal{F}$ *that is a strict subset of* A*, then for each* $i \in A \backslash A'$ *we have* $A \backslash \{i\} \in \mathcal{F}$ *by comprehensiveness, and therefore* $s_i^A\left(b_i, b_{N\backslash A}\right) > 0$ *by perfect feasibility checking, hence the algorithm cannot stop at* A*.)*

In some settings, however, perfect feasibility checking may be computationally impossible, and imperfect checking must be used instead. For example, in the FCC's spectrum-clearing problem, a given set A *of bidders is in* $\mathcal{F}$ *when there exists an assignment of the rejected bidders* $N \backslash A$ *to available channels that satisfies all interference constraints (as formalized in (3) below), so checking that requires solving an NP-hard problem. When a feasibility checker has a limited time to run on a computationally challenging problem, it may generate one of the following three outputs: (i) a proof that* $A \backslash \{i\} \in \mathcal{F}$ *(which consists of a feasible assignment of bidders* $(N \backslash A) \cup \{i\}$ *to channels), (ii) a proof that* $A \backslash \{i\} \notin \mathcal{F}$ *(which could be a long logical argument), or (iii) a time-out, which means that the algorithm had to be terminated before producing either (i) or (ii). In case (i), we set* $s_i^A\left(b_i, b_{N\backslash A}\right) > 0$*,*

[11] E.g., in the example of Section 2, $\mathcal{F} = \{\{1,2\}, \{3\}, \{1,3\}, \{2,3\}, \{1,2,3\}\}$.

while in cases (ii) and (iii), we need to set $s_i^A\left(b_i, b_{N\setminus A}\right) = 0$ *in order to guarantee that the algorithm always yields a feasible assignment. Without perfect feasibility checking, the outcome may not be non-wasteful.*[12]

Example 2 (Optimization with Matroid Constraints) *Suppose that the goal is to find an "efficient" (i.e. social cost-minimizing) set of winning bids subject to the feasibility constraint in Example 1:*

$$\alpha\left(b\right) \in \arg\min_{A\in\mathcal{F}} \textstyle\sum_{i\in A} b_i. \tag{2}$$

Consider the "standard" DA algorithm that does perfect feasibility checking and scores the bids that are feasible to reject by their bid amounts, i.e., $s_i^A\left(b_i, b_{N\setminus A}\right) = \begin{cases} b_i & \textit{when } A\setminus\{i\}\in\mathcal{F}, \\ 0 & \textit{otherwise,} \end{cases}$ *and assume that there are no ties (i.e.,* $b_i \neq b_j$ *for* $i \neq j$ *for each* $b \in B$*). By a classical result in matroid theory (see Oxley (1992)), the standard DA algorithm computes an efficient allocation rule* α *in this setting if and only if the set family* $\{R \subseteq N : N\setminus R \in \mathcal{F}\}$ *– i.e., the feasible sets of* rejected *bids – form the independent sets of a matroid on the ground set* N*.*[13]

[12] Note that if the set $\mathcal{F}$ is comprehensive, then once the infeasibility of adding a bidder has been established, the bidder will remain infeasible to add in all the subsequent rounds, hence there is no point in running the checker again for the same bidder. In contrast, if the feasibility checker times out on a given bidder, it may be run again in the same station in a subsequent round in which more computational time is available (e.g., overnight). If the checker does find a feasible way to add this bidder, this will permit us to either in reject the bidder or reduce his payment.

[13] This is closely related to the analysis of Bikchandani et al. (2011), who focus on "selling"' auctions in which the family of sets of bids that could be feasibly *accepted* is a matroid. For these settings, the DA heuristic achieving efficiency is the "greedy worst-out" heuristic: reject the lowest bid that could be rejected while the set of accepted bids still forms a spanning set of the matroid. (Note that in contrast to the "standard" DA procurement auction, the provisional allocation in the selling auction remains infeasible until the end of the auction.) By the argument in Section 6 below, this DA heuristic corresponds to the ascending-price clock auction of Bikchandani et al. (2011), which offers the same ascending price to all the active bidders that could still be rejected. While both the "greedy best-in" and the "greedy worst-out" algorithms can be used to compute the efficient allocation rule in both the procurement and selling matroid auctions, only one

Intuitively, the matroid property in Example 2 involves a notion of "one-for-one substitution"' between bidders. (In particular, it implies that all the maximal feasible sets of rejected bids – the "bases" of the matroid – have the same cardinality.) This does not hold, even approximately, in the FCC setting, which involves a trade-off between acquiring a larger number of stations with smaller coverage areas and a smaller number of stations with larger coverage areas, as illustrated in Section 2. In such cases, instead of the standard DA algorithm it is desirable to use scoring that assigns greater scores to larger stations, given the same bid values. This is illustrated in the following stylized example:

Example 3 (Scoring for the Knapsack Problem) *Suppose that the family of feasible sets takes the form*

$$\mathcal{F} = \left\{ A \subseteq N : \textstyle\sum_{i \in N \setminus A} w_i \leq 1 \right\}.$$

The problem of maximizing total surplus subject to this constraint is known as the "knapsack problem." Interpreting $w_i > 0$ as the "volume"' of bidder i, the problem (2) is to minimize the total value of the accepted bids subject to a constraint on their total volume.[14] *While this optimization problem is NP-hard, it can be approximated with the famous heuristic of Dantzig (1957), which corresponds to the DA algorithm with the "per-volume" scoring* $s_i^A\left(b_i, b_{N\setminus A}\right) = \begin{cases} b_i/w_i & \textit{when } A\setminus\{i\} \in \mathcal{F}, \\ 0 & \textit{otherwise.} \end{cases}$.

In the FCC setting, keeping total procurement cost low is also an important goal. A cost-minimization objective can be easily handled using the "virtual value" approach of Myerson (1981).

of them is a DA algorithm: it is the best-in algorithm for rejecting bids in the procurement auction, and the worst-out algorithm for rejecting bids in the selling auction.

[14] Duetting, Gkatzelis and Roughgarden (2014) consider instead the "selling" problem in which the *accepted* bids must fit into a knapsack, and examine the approximation power of DA algorithms for this problem.

Example 4 (Scoring for Expected Cost Minimization) *Suppose that bidders' values v_i are independently drawn from distributions $F_i(v) = \Pr\{v_i \leq v\}$, and that $B_i = [0, \bar{b}_i]$ for each i (so that bidding exact values is possible). Then, following the logic of Myerson (1981), the expected cost of a threshold auction that implements allocation rule α can be expressed as $\mathbb{E}\left[\sum_{i\in\alpha(v)} \gamma_i(v_i)\right]$, where $\gamma_i(v_i) = (v_i + F_i(v_i)/F_i'(v_i))$ (bidder i's "virtual cost function"). Assume the virtual cost functions are strictly increasing. Then, if we are given a DA algorithm with scoring rule s that exactly or approximately minimizes the expected social cost subject to feasibility constraints as in the above examples, it can be modified to yield a DA threshold auction that exactly or approximately minimizes the total expected cost of procurement subject to the constraints. The modified auction uses the scoring rule $\hat{s}_i^A\left(b_i, b_{N\setminus A}\right) = s_i^A\left(\gamma_i(b_i), \left(\gamma_j(b_j)\right)_{j\in N\setminus A}\right)$.*

The previous examples have used "fixed" scoring, which conditions only on the bidder's identity and his bid value. The following two examples demonstrate two reasons why it may be useful to use "adaptive" scoring, that is, to condition on the set of rejected bids and their bid values.

Example 5 (Adaptive Scoring for a Budget Constraint) *Suppose that the designer faces the "budget constraint" that the total payment to the winners cannot exceed $R(A)$ when the set of accepted bids is $A \subseteq N$. In order to find a maximal set of winners satisfying the budget constraint, the algorithm needs to continue whenever the total threshold prices to be paid to the winners exceeds $R(A)$. If the budget constraint is the only limit on procurement, then the algorithm stops as soon as that limit is satisfied. For a simple example, if scoring were based on increasing functions $\sigma_i : B_i \to \mathbb{R}_{++}$ (depending only on the i's bid), then the threshold price that would have to be paid to any currently active bidder $i \in A$ if the algorithm were to stop right away is:*

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

In this case, the "stopping rule" described above can be implemented via the scoring functions

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

These kinds of algorithms may be called "revenue-sharing"' algorithms, and their mirror images for the "selling"' auction are the "cost-sharing"' algorithms of Moulin (1999) and Mehta et al. (2009). (While this formulation permits the auction to generate a positive net revenue, it is possible to modify it to require exact budget balance.) An equivalent clock auction has been considered by Ensthaler and Giebe (2014) for the case where $R(A)$ is a constant.

It is also possible to accommodate both budget constraints and other feasibility constraints. For example, consider McAfee's (1992) double-auction algorithm for a homogeneous-good market with unit buyers and unit sellers, in which the feasibility constraint dictates that the number of the accepted buying bids ("demand") equal the number of the accepted selling bids ("supply"). McAfee's algorithm can be viewed as a DA threshold auction that rejects either the highest-cost selling bid or the lowest-value buying bid in each round, making sure that demand is always be within 1 of supply, and that stops as soon as the twin conditions are met: (i) demand equals supply and (ii) the net threshold payment is nonnegative. (This mechanism is equivalent to a clock auction in which all buyers face the same ascending price and all sellers face the same descending price.) The FCC's "incentive auction" is similarly a double auction for spectrum sellers (TV broadcasters) and spectrum buyers (mobile broadband companies) that is constrained to generate a certain amount of net revenue, but subject to the added complication that buyers demand and sellers supply different kinds of differentiated goods and the feasible combinations of accepted bids are quite complicated. However, FCC's problem also admits a DA double auction that is similar to McAfee's double auction, which sets a sequence of n possible spectrum targets for the *number of channels to clear, and starts by trying to clear the largest number, reducing the target whenever the revenue goals cannot be achieved with the current spectrum clearing target.*[15]

Example 6 (Adaptive Scoring for Yardstick Competition) *Suppose that the auctioneer cares about the expected total profit, that her gross profit for acquiring each bidder is* π*, and that there are no other constraints on the feasible allocations. Further suppose that bidders' values are random, i.i.d. draws from a distribution that is unknown to the auctioneer. In this symmetric setting we can focus without loss on auctions that treat bidders symmetrically. The simplest DA algorithms simply set a reserve price* p^* *and buy from all bidders who quote a lower price. This can be implemented as a DA auction with* $s_i^A(b_i) = \max\{v_i - p^*, 0\}$*) to yields an expected profit on each bidder* i *of* $(\pi - p^*)\Pr\{v_i \le p^*\}$*, where the probability is calculated based on the auctioneer's prior. However, it is evident that the auctioneer can do better in expectation by conditioning the price quoted to each still-active bidder on the information implied by the previously rejected bids. Letting* $p_A^*\left(b_{N\setminus A}\right) = \arg\max_p (\pi - p)\Pr\left\{v_i \le p | b_{N\setminus A}\right\}$ *be the optimal monopsony price for the posterior distribution of values given the rejected bids, the better DA auction would set scores according to* $s_i^A\left(b_i, b_{N\setminus A}\right) = \max\left\{v_i - p_A^*\left(b_{N\setminus A}\right), 0\right\}$.[16]

[15] Duetting, Roughgarden, and Talgam-Cohen (2014) consider the approximation power of balanced-budget DA double auctions for settings in which buyers and sellers must be matched one-to-one subject to some constraints.

[16] Auctions that are even more profitable are possible. The expected profit-maximizing threshold auction examined in Segal (2003) implements the allocation rule $\alpha(b) = \{i \in N : b_i \le p^*(b_{-i})\}$, that is, it offers each bidder the optimal price given the information inferred from all the other bids. However, this allocation rule is typically not implementable with a DA algorithm, because in it the winning bidders could be “bossy”' – affect the allocation while still winning. In contrast, in a DA algorithm winning bidders are always non-bossy.

5.1 Example: Near-Optimality of DA Algorithms

For computationally challenging problems in which full optimization is impossible, one may ask where there exists some DA algorithm achieve reasonably good performance. Partly, this question should be addressed by simulations. In this subsection, we examine a setting in which this question can be studied theoretically and which is similar to the FCC repacking problem.

In this setting, the bidders are television stations who bid to relinquish their broadcast rights. Broadcast channels must be assigned to the stations whose bids are rejected in a way that satisfies non-interference constraints. In this simplified model, those constraints are represented by an interference graph Z: a set of two-elements subsets of 2^N ("edges"). We interpret $\{i,j\} \in Z$ to mean that stations i and j cannot both be assigned to the same channel without causing unacceptable interference. Letting $\{1,...,n\}$ denote the set of channels left after the auction. Then, the feasible sets of accepted bids for the FCC's repacking problem can be written as

$$\mathcal{F} = \{A \subseteq N : (\exists c : N\backslash A \to \{1,...,n\})(\forall i,j \in N\backslash A)(\{i,j\} \in Z \Rightarrow c(i) \neq c(j)\}. \tag{3}$$

Proposition 2 *Suppose there exists an ordered partition of the set N of stations into m disjoint sets $N_1,...,N_m$ such that*

(i) for each $k = 1,...,m$ and each $i,j \in N_k$, $\{i,j\} \in Z$ (that is, each N_k is a "clique"),

(ii) there exists some $d < n$ such that for each $k = 1,...,m$ and each $S \subseteq N_k$ satisfying $|S| \leq n-d$, we have

$$|S| + |\cap_{i \in S} \cup_{l<k} \{j \in N_l : \{i,j\} \in Z\}| \leq n.$$

Then there exists a DA algorithm that, for any set of bids, first rejects the most valuable stations in each partition element N_k, iterates in that way *as long as that is feasible, and then continues in any way. For every such algorithm, the value of stations assigned is at least a fraction* $1 - d/n$ *of the optimal value.*

Intuitively, the Proposition (proven in Appendix C) applies to the setting in which the stations can be partitioned into m "metropolitan areas" in such a way that (i) no two television stations in the same area can be assigned to the same channel and (ii) the cross-area constraints are limited in the sense that if we have a set S of no more than $n - d$ stations in one metropolitan area, there are no more than $n - |S|$ stations in lower-indexed areas that have interference constraints with *all* the stations in S.[17] Using an argument based on Hall's Marriage Theorem, condition (ii) ensures that it is possible to select any arbitrary $n-d$ stations in each area independently of each other and still be able to find a feasible assignment of these stations to channels. Since the optimal value is bounded above by assigning the n most valuable stations in each area (which would be feasible if there were no inter-area constraints), the worst-case fraction of efficiency loss is bounded above by $(n - d)/n$.

We make several observations.

1. It is possible to satisfy condition (ii) while having several times more inter-area constraints than within-area constraints. To illustrate, suppose that all stations are arranged from north to south on a line and that each station interferes with its $n-1$ closest neighbors to the north as well as its $n-1$ closest neighbors to the south. Suppose that each successive group of n stations is described as a metropolitan area. Then, it is possible to assign all stations ($d = 0$) to channels without creating interference just by rotating through the n channel numbers. In this example, there are just $x = n(n-1)/2$ constraints among stations within

[17] While we assume that the partition is totally ordered for notational simplicity, the proposition can also be extended to cases in which partition elements form an ordered acyclic graph, by interpreting $<$ and $\leq$ as referring to a precedence relation.

an area but $2x$ constraints between those stations and ones in the next lower indexed area and another $2x$ constraints involving stations in the next higher indexed area.

2. In general, there may be many ways to partition stations into cliques, and many ways to order any given partition. The Proposition formally applies to each partition, but the number d and therefore the approximation guarantee will vary depending on which partition is selected and how it is ordered.

3. The worst-case bound applies over all possible station values and incorporates both a conservatively high estimate of the optimum and a conservatively low estimate of the DA algorithm performance. To approach the worst-case bound, the optimum must be similar to assigning the n most valuable stations in each area, and the DA algorithm must be unable to assign any other stations after the $n - d$ most valuable stations have been assigned in each area.

4. The standard DA algorithm discussed in Example 2 is not among the ones described in the Proposition, and does not achieve the same performance guarantee. For example, suppose that there is some central area – area 1 – such that the stations in any area are linked to all other stations in that area and to the stations in area 1, but not to any other stations. Suppose that there are 2 channels available. Then the DA algorithms described in the proposition assign the single most valuable station in each area, thus achieving at least half of the optimal value. On the other hand, the standard DA algorithm could in this case achieve as little as $1/(m-1)$ of the optimal value: This could happen if the two most valuable stations happen to be in area 1, in which the standard DA algorithm assigns just those two stations and no others. Thus, the example in this section demonstrates how it may be possible to design a DA algorithm to improve upon the standard

DA algorithm by taking advantage of known properties of the feasible set. In applications like the FCC auction, in which interference graph is known before the auction, it may be possible to apply a variety of analytical tools and simulations to find a DA algorithm that performs much better than the standard one.

6 Clock Auctions

Informally, a ("descending") clock auction is a dynamic mechanism that presents a declining sequence of prices to each bidder, with each presentation followed by a decision period in which any bidder whose price has been strictly reduced decides whether to exit or continue. Bidders that have not exited are called "active". Bidders who remain active when their prices are reduced are said, informally, to "accept" the lower price. When the auction ends, the active bidders become the winners and are paid their last (lowest) accepted prices. Different clock auctions are distinguished by their pricing functions, which determine the sequence of prices presented.[18]

To formalize the intuitive description, we restrict attention to finite clock auctions with discrete periods.[19] The active bidders in period t are denoted by $A_t \in 2^N$. A period-t history consists of the sets of active bidders in all periods up to period t: $A^t = (A_1, ..., A_t)$ such that $A_t \subseteq ... \subseteq A_1 = N$. Let H denote the set of all such histories. A descending clock auction is 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.)

[18]The described descending clock auctions for the procurement setting are the mirror image of the ascending clock auctions for the selling setting, which in turn generalize the classic English auction for selling a single item.

[19]We say that the clock auction is *finite* if there exists some T such the auction always stops by period T. The restriction to finite auctions avoids familiar technical difficulties, such as those associated with describing continuous-time auctions and with defining dominance solvability for games with infinite strategy spaces.

The clock auction initializes $A_1 = N$. In each period $t \geq 1$, given history A^t, a profile of prices $p(A^t)$ is offered to the bidders. If $p(A^t) = p(A^{t-1})$, the auction stops; 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 $i \in A_t$ and $p_i(A^t) < p_i(A^{t-1})$, then i may choose to refuse the new price and exit the set of active bidders.[20] 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 $A^{t+1} = (A^t, A_{t+1})$.

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 that the bidder should 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 DA threshold auctions, meaning that the mapping from the profile of bids or cutoffs to allocations and prices are the same for both auctions.[21]

Proposition 3 *For every DA threshold auction with finite bid spaces, there exists an equivalent finite clock auction in which bidders are restricted to cutoff strategies.*

[20] In a variant of the auction, all active bidders $i \in A_t$ may choose whether to exit. Although the results below are the same for the auction in the main text and this variant, the version in the main text is preferred because, when there is a feasibility constraint as in 1, it ensures that the clock auction always yields a feasible outcome.

[21] If we were to implement a deferred-acceptance threshold auction as a multi-round procedure in which some information is disclosed to active bidders between rounds and the active bidders are allowed to improve their bids, then the resulting mechanisms would be the "survival auctions" proposed by Fujishima et al. (1999). These auctions are strategically equivalent to clock auctions without the restriction to cutoff strategies.

Proof. Given bid spaces $B_1, .., B_N$, for each $v \in \mathbb{R}$, let $v^+ = \min\{b_i \in B_i : b_i > v\}$ 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 DA threshold auction with scoring rule s, we construct an equivalent clock auction in which the price to every highest-scoring active bidder is reduced by the minimal amount, while the prices for all other bidders are left unchanged:

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

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

Then equivalence is easy to see: First, for every history of the clock auction, when bidders use cutoff strategies, the next set of bidders to exit in the clock auction is the set of bidders who have the maximum scores among the set of active bidders given bidders' cutoffs, and iterating this argument establishes that the final set of winners is the same in both auctions. Second, for each bidder who has not exited by the auction's end and thus became a winner, his final clock price is the highest cutoff it could have used to be winning – i.e., his threshold price. ■

Proposition 4 *For every finite clock auction in which bidders are restricted to cutoff strategies, there exists an equivalent DA threshold auction with finite bid spaces.*

Proof. Given a finite clock auction p, we construct bid spaces and a scoring rule to create an equivalent DA threshold auction. 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 algorithm 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 exited 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 algorithm 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. ■

7 Group Strategy-proofness

Definition 4 *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 5 *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 1 *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 5 *Every finite clock auction (with any information disclosure) is weakly group strategy-proof.*

Proof. Every truthful bidder has a payoff of at least zero. Consider the first stage of clock auction affected by a group deviation. A bidder cannot benefit by a deviation that results in it losing and getting a deviation payoff of zero, so the first deviation must be made by a bidder who chooses *not* to exit at a price equal to or below his truthful value, and who becomes winning. That bidder's price, however, must then be below his truthful value, so that bidder loses from the group deviation. ■

The preceding argument is independent of the information policy in the auction, but it applies in particular when bidders are restricted to play cutoff strategies. Combining the two previous propositions, we get

Corollary 1 *Any DA threshold auction is weakly group strategy-proof.*[22]

[22] One can also establish a version of this result directly without restricting attention to finite bid spaces.

8 Pay-as-Bid: Full-info equivalence

Recall that for any bid space B and allocation rule α, the rounded up values are $v^+ = \min\{b_i \in B_i : b_i > v\}$ and 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 6 *For every paid-as-bid DA auction with bid sets B_i and every value profile v with $v_i < \max B_i$ for each $i \in N$, there is a complete-information Nash equilibrium bid profile b with $b_i = \max\left\{v_i^+, p_i\left(v_{-i}^+\right)\right\}$ for each $i \in N$. The corresponding equilibrium allocation is $\alpha(b) = \alpha(v^+)$.*

Proof. Let $A = \alpha(v^+)$, and note that the definition of threshold prices implies that

$$b_i = \max\left\{v_i^+, p_i\left(v_{-i}^+\right)\right\} = \begin{cases} p_i\left(v_{-i}^+\right) & \text{for } i \in A, \\ v_i^+ & \text{for } i \in N \backslash A. \end{cases}$$

Since changing accepted bids so that they are still accepted does not alter the outcome of the DA algorithm, we have

$$\begin{aligned} \alpha\left(v^+\right) &= \alpha\left(p_A\left(v^+\right), v_{N\backslash A}^+\right) = \alpha(b), \text{ and} \\ p_i\left(v_{-i}^+\right) &= p_i\left(p_{A\backslash i}\left(v^+\right), v_{N\backslash A}^+\right) = p_i(b_{-i}) \text{ for each } i \in A. \end{aligned}$$

Now, we verify that the bid profile b is a Nash equilibrium. Every bidder $i \in A$ is winning by bidding $b_i = p_i\left(v_{-i}^+\right) = p_i(b_{-i}) \geq v_i^+$, and since any bid above $p_i(b_{-i})$ would make him lose, he has no profitable deviation. Every bidder $i \in N\backslash A$ is losing with his bid of $b_i = v_i^+$, and since he could only win by lowering his bid, he has no profitable deviation. ■

Proposition 6 only describes one Nash equilibrium outcome among many. In the remainder of this section, we show how the described outcome may be selected as the unique prediction by eliminating dominated strategies (either iteratively, or in a single round and then considering Nash equilibria in the remaining strategies). However, in order to ensure uniqueness, we need to eliminate examples like the following one (which could arise as a particular instance of yardstick competition described in Example 6):

Example 7 *Let* $N = 2$*,* $B_1 = \{1, 3\}$*,* $B_2 = \{2, 4\}$*,* $\alpha(b) = \{1\}$ *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. There are two Nash equilibrium profiles in undominated strategies: they are* $(3, 4)$ *(in which bidder 1 wins) and* $(3, 2)$ *(in which there is no winner). Both strategy profiles also survive iterated deletion of dominated strategies.*

Intuitively, what leads to multiplicity of outcomes in the example is that neither iterated dominance nor undominated Nash equilibrium nails down the behavior of the losing bidder 2, and this behavior in turn affects the allocation of bidder 1. In order to rule out such situations, we restrict attention to assignment rules that are non-bossy:

Definition 6 *Allocation rule* α *is* non-bossy *if for any* $i \in N$*,* $b \in B$ *and* $b_i' \in B_i$*,* $\alpha(b_i', b_{-i}) \cap \{i\} = \alpha(b) \cap \{i\}$ *implies* $\alpha(b_i', b_{-i}) = \alpha(b)$*.*

Non-bossiness means that a bidder cannot affect others' allocations without changing his own allocation. In a DA algorithm, a winner who changes its bid without changing its winning status (that is, an agent $i \in \alpha(b_i, b_i) \cap \alpha(b_i', b_i)$) can never affect others' winning status, but because bidder's scores can depend on losing bids, a change by a loser to a different losing bid ($i \notin \alpha(b_i, b_i) \cup \alpha(b_i', b_i)$) could affect the set of winners, like it does in the above example. The non-bossiness condition rules this out. Examples of DA algorithms that satisfy this condition will be given below.

Definition 7 *An auction is* dominance-solvable *in state* v *if under full information, there exists a unique payoff profile that remains after iterated deletion of (weakly) dominated strategies, regardless of the order of elimination.*

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

Intuitively, iterated deletion of weakly dominated strategies closely resembles a deferred-acceptance procedure, because it works by iterated rejections using a myopic criterion and finally accepting all strategies that are not rejected. We find that, for paid-as-bid auctions, iterated elimination of weakly dominated bids is closely related to the iterated deletions of always-losing bids that characterize DA auctions, and that this implies a similarly close connection between dominance-solvable auctions and DA auctions:

Proposition 7 *Consider a paid-as-bid auction with a monotonic, non-bossy allocation rule α and finite bid spaces B. Say that a value profile v is "generic" if $v_i \in [0, \bar{v}_i] \setminus B_i$ for each i.*

(i) The auction is pure-strategy dominance-solvable for all generic value profiles if and only if α can be implemented with a DA algorithm.

(ii) If α can be implemented with a DA algorithm, then for every generic value profile, the unique payoff profile surviving iterated deletion of dominated strategies is also the unique payoff profile associated with any (pure or mixed) Nash equilibrium in undominated strategies.

(iii) If α can be implemented with a DA algorithm, then one strategy profile that survives iterated deletion of dominated strategies and is a Nash equilibrium in undominated strategies is the one described in Proposition 6.

The equivalence of dominance solvability and implementation using a DA algorithm – equivalently, using the results of Section 6, a clock-auction algorithm – is constructive and intuitive. Indeed, a clock auction decrements the price to a bidder when the bidder could no longer win at that price, given the prices already accepted by the other bidders. Similarly, in a round of iterated deletion of dominated strategies in a paid-as-bid auction, a bid may be dominated because it always loses against all remaining possible bid profiles of the other bidders (when a lower bid, still above the bidder's value, might sometimes be winning). The extra work in the proof arises because, in a paid-as-bid auction, bids may also be dominated in other ways: A bid is also dominated when it is below the bidder's value, or when there exists a higher bid that wins against exactly the same opposing bid profiles. In a non-bossy auction, by a result of Marx and Swinkels (1997), dominated bids can be eliminated in any order without affecting the final outcome of the auction. The proof of the proposition, given in Appendix D, distinguishes the reasons for domination to identify steps that correspond to price reductions in a clock auction.

Now we describe two classes of non-bossy allocation rules. One class contains efficient allocation rules (recall from Example 2 that in some cases such rules can be computed by DA heuristics):

Lemma 1 *Consider an efficient allocation rule α given by (2) and assume that there are no ties (i.e.* $\arg\min$ *is always single-valued). Then the allocation rule is non-bossy.*

Proof. (2) implies that for all $i \in N$, $b_i, b_i' \in B_i$, $b_{-i} \in B_{-i}$,

$$\begin{aligned} i &\notin \alpha\left(b_i, b_{-i}\right) \cup \alpha\left(b_i', b_{-i}\right) \Longrightarrow \alpha\left(b_i, b_{-i}\right) = \arg\min_{A \in \mathcal{F}: i \notin A} \textstyle\sum_{j \in A} b_j = \alpha\left(b_i', b_{-i}\right), \\ i &\in \alpha\left(b_i, b_{-i}\right) \cap \alpha\left(b_i', b_{-i}\right) \Longrightarrow \alpha\left(b_i, b_{-i}\right) = \arg\min_{A \in \mathcal{F}: i \in A} \textstyle\sum_{j \in A \setminus \{i\}} b_j = \alpha\left(b_i', b_{-i}\right). \end{aligned}$$

■

Another class contains allocation rules computed by DA heuristics with fixed scoring and perfect feasibility checking:

Lemma 2 *Suppose that $\mathcal{F} \subseteq 2^N$ is a comprehensive family of feasible sets (see Example 1), and that scoring is given by* $s_i^A\left(b_i, b_{N\setminus A}\right) = \begin{cases} \sigma_i\left(b_i\right) & \text{if } A\setminus\{i\} \in \mathcal{F}, \\ 0 & \text{otherwise,} \end{cases}$ *where $\sigma_i : B_i \to \mathbb{R}_{++}$ are strictly increasing functions that have no ties (so feasibility is always maintained). The allocation rule computed by the resulting DA heuristic is non-bossy.*

Proof. As observed above, every DA procedure satisfies non-bossiness for the accepted bids. To check that condition for the rejected bids, too, we show that if given bid profile b agent i's bid b_i is rejected in some round t, then replacing his bid with some $b_i' > b_i$ does not affect the final outcome of the algorithm. Note that it suffices to check situations in which the replacement results in b_i' rejected *prior to* round t: Indeed, otherwise b_i' will be rejected in round t and the replacement will not affect the behavior of the algorithm. Suppose first that the replacement results in b_i' being rejected in round $t-1$, and thus does not affect the behavior of the algorithm prior to round $t-1$. Letting A_{t-1} be the set of accepted bids in round $t-1$, and letting agent j be the agent whose bid is rejected in round $t-1$, we must have $A_{t-1} \setminus \{i, j\} \in \mathcal{F}$, and

$$\max_{k \in A_{t-1} \setminus \{i,j\} : A_{t-1} \setminus \{j,k\} \in \mathcal{F}} \sigma_k (b_k) < \sigma_i (b_i) < \sigma_j (b_j) .$$

Using this and the comprehensiveness of $\mathcal{F}$, we obtain

$$\max_{k \in A_{t-1} \setminus \{i,j\} : A_{t-1} \setminus \{i,j,k\} \in \mathcal{F}} \sigma_k (b_k) \leq \max_{k \in A_{t-1} \setminus \{i,j\} : A_{t-1} \setminus \{j,k\} \in \mathcal{F}} \sigma_k (b_k) < \sigma_j (b_j) ,$$

which implies that bid b_j must be rejected in round t. Then after round t the algorithm will be unaffected by the replacement of b_i with b_i'. Iterating this argument, we see by induction on τ that any increase in agent i's bid, resulting in it being rejected in some round $t-\tau$, will not affect the allocation.
■

9 Conclusion

The analysis and results reported in this paper were developed in response to questions arising from an engagement to advise the FCC on the design of its "incentive auction." The possibility that some deferred acceptance algorithm could achieve nearly optimal results is of obvious importance for that application: it inspired simulations to identify the DA algorithms that work best for the FCC's particular interference graph. The ability to implement any deferred acceptance threshold auction with a clock auction has been important both because such auctions may be easier for bidders than comparable sealed-bid auctions, and because a clock auction remains weakly group strategy-proof even if bidders do not trust the auctioneer's computations. Strategy-proofness is also important because it reduces the cost of participation, especially for small local broadcasters whose participation is needed for a successful incentive auction. Our outcome equivalence results suggest that strategy-proofness can be introduced without raising acquisition costs compared to a paid-as-bid rule. Since there can be many Nash equilibria, the credibility of the outcome equivalence result is enhanced by the related outcome uniqueness result. In the incentive auction, revenues for the forward auction portion must be sufficient to pay the costs the broadcasters incur in moving to new broadcast channels, as well as meeting certain other gross and net revenue goals, so the possibility of including a cost target in the scoring rule is necessary to make the whole design possible. Including yardstick competition allows the FCC to set maximum prices for broadcasters in regions with little competition based on bids in other, more competitive regions.

One important limitation of our analysis is its focus on single-minded bidders. In practice, bidders are often interested in selling or buying various goods or bundles and have different values for those packages. For example, in the FCC's "reverse auction," in addition to selling their UHF stations to go off-air, bidders may also be permitted to bid to switch to a less-congested lower-frequency VHF band. Also, some broadcasters own multiple stations and may interested in selling different subsets of their stations.

Deferred-acceptance auctions and clock auctions have also been used and studied for multidimensional bidders. For some special settings, DA auctions have been proposed that are strategy-proof (e.g. Kelso and Crawford 1987, Ausubel 2004, Bikhchandani et al. 2011), but other kinds of proposed auctions are not. For example, in the simultaneous multiple-item ascending auctions studied by Milgrom (2000) and Gul and Stachetti (2000), bidders who may be willing to purchase more than one item at a time have an incentive to engage in monopsonistic "demand reduction." Such incentives may be minimized in "large-market" settings in which competition limits the power of each bidder to affect prices, ensuring approximate strategy-proofness. Another approach is to design exactly strategy-proof clock auctions that do not guarantee exact efficiency (see, e.g., Bartal et al. (2003). Further examination of DA auctions for multidimensional bidders remains an important direction for future research.

Roth (2002) has observed that "Market design involves a responsibility for detail, a need to deal with all of a market's complications, not just its principal features." Over the past two decades, variants of the original DA algorithm have had remarkable success in accommodating the diverse details and complications of a wide set of matching market design problems. In this paper, we extend that success to new class of auction design problems, connect DA auctions to encompass clock auctions, and reaffirm the deferred-acceptance idea as the basis for many of the most successful mechanisms of market design.

References

[1] Aardal, K., S. P. M. van Hoesel, A. M. C. A. Koster, C. Mannino, and A. Sassano (2007), "Models and solution techniques for frequency assignment problems," *Annals of Operations Research* 153(1), 79-129.

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

[3] Ausubel, L.M. and P. Milgrom (2002), "Ascending Auctions with Package Bidding, " *Frontiers of Theoretical Economics*, 1 (1), 1452–1475.

[4] Ausubel, L.M. and P. Milgrom (2006), "The Lovely But Lonely Vickrey Auction," in *Combinatorial Auctions*, P. Cramton, Y. Shoham, and R. Steinberg (eds), Cambridge, MA:MIT Press, 2006, chapter 1.

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

[6] 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.

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

[8] 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.

[9] Dantzig, G. (1957), "Discrete-variable extremum problems," *Operations Research* 5, 266–277

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

[11] Duetting, P., V. Gkatzelis, and T. Roughgarden (2014), "The Performance of Deferred-Acceptance Auctions," *ACM Conference on Electronic Commerce*

[12] Duetting, P., T. Roughgarden, and I. Talgam-Cohen (2014), "Modularity and Greed in Double Auctions," *ACM Conference on Electronic Commerce*

[13] 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.

[14] Ensthaler, L., and T. Giebe (2014), "A dynamic auction for multi-object procurement under a hard budget constraint," *Research Policy* 143, 179-189

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

[16] Gale, D., and L.S. Shapley (1962), "College Admissions and the Stability of Marriage," *American Mathematical Monthly* 69, 9–14

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

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

[19] Hatfield, J. and F. Kojima (2009), "Group Incentive Compatibility for Matching with Contracts," *Games and Economic Behavior* 67: 745-749.

[20] Hatfield, J., and P. Milgrom (2005), "Matching with Contracts," *American Economic Review* 95(4), 913-935.

[21] 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.

[22] 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.

[23] 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.

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

[25] Klemperer, P. (2010). "The Product-Mix Auction: A New Auction Design for Differentiated Goods," *Journal of the European Economic Association* 8(2-3), 526-536.

[26] Kojima, F., P. Pathak and A. Roth (2013), "Matching with Couples: Stability and Incentives in Large Markets," *Quarterly Journal of Economics.*

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

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

[29] McAfee, R.P., (1992), "A Dominant Strategy Double Auction," *Journal of Economic Theory* 56(2), 434–450

[30] McMillan, J. (1994), "Selling Spectrum Rights," *Journal of Economics Perspectives*, 8(3), 145-162.

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

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

[33] Milgrom, P. (2004), *Putting Auction Theory to Work*, Cambridge University Press.

[34] Milgrom, P. (2009), "Assignment Messages and Exchanges," *American Economic Journal: Microeconomics*, 1(2), 95-113.

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

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

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

[38] Myerson, R. (1981), "Optimal Auction Design," *Mathematics of Operations Research 6, 58-73.*

[39] Nisan, N., and A., Ronen (1999), "Algorithmic Mechanism Design," *Proceedings of the thirty-first annual ACM Symposium on Theory of Computing (STOC)*, 129-140.

[40] Oxley, James G. (1992). *Matroid Theory.* Oxford University Press, Oxford, GB.

[41] Roth, Alvin E. and Elliot Peranson (1999). "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design". *American Economic Review* 89 (4): 748–780.

[42] Segal, I. (2003), "An Optimal Pricing Mechanism with Unknown Demand," *American Economic Review* 93(3), 509-529.

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

[44] Singer, Y. (2010), "Budget Feasible Mechanisms," *51st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 765 - 774.

[45] Weber, R.J. (1997), "Making More from Less: Strategic Demand Reduction in the FCC Spectrum Auctions," Journal of Economics & Management Strategy, 6 (3), 529–548.

10 Appendix A: Greedy-Acceptance Auctions

In this appendix, we illustrate the greedy-acceptance auctions of Lehmann, O'Callaghan and Shoham (2002). To be consistent with their setting, we consider a three-bidder "forward" (selling) auction (although it is also possible to construct a corresponding reverse auction) in which it is feasible to satisfy either both bidders 1 and 2 or bidder 3. (Say, bidder 3 desires a bundle of two objects while bidders 1 and 2 each desire a single object from the bundle.) For a simple example, the bidders' scores could be defined as their bids. The auction iterates accepting the highest bid that is still feasible to accept. The winners are paid their threshold prices, i.e., the minimal bids that would have been accepted.

Consider the case in which both bidders 1 and 2 bid above bidder 3's bid. In this case they both win and pay zero, since each of them would have still won by bidding zero, letting the other bid be accepted in the first round, which makes bid 3 infeasible to accept. This implies that the auction is not weakly group strategyproof: when the true values of bidders 1,2 are below b_3 but strictly positive, they would both strictly benefit from both bidding above b_3, so that they both win and pay zero. Also, the auction's revenue in this case is zero, while any full-information Nash equilibrium of the corresponding paid-as-bid auction could not sell to bidders 1 and 2 at a total price below bidder 3's value, since otherwise bidder 3 could have profitably deviated to win the auction. Finally, this allocation rule cannot be implemented with DA auction (which is equivalent to an ascending-price clock auction in the forward-auction setting), since the allocation is completely determined by the single highest bid while the first step of the DA algorithm is determined by the single worst (lowest) bid according to some criterion. Nor could it be implemented with a descending-clock ("Dutch") auction, since in that auction, when bidder 3 buys first, the allocation is determined but the winner's threshold payment is not yet determined.[23]

[23] A traditional purpose of a clock auction is to economize on information transmission or

11 Appendix B: Substitutable DA rules

We say that a clock auction *implements allocation rule* α if $\{p_i(h) : h \in H\} = B_i$ for each i and the auction generates $\alpha(b)$ when bidders use cutoff strategies with cutoffs $b \in B$. Recall that allocation rule α is *monotonic* if $i \in \alpha(b)$ and $b'_i < b_i$ implies $i \in \alpha(b'_i, b_{-i})$. Say that allocation rule α has *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 8 *With finite bid spaces, any monotonic allocation rule with substitutes whose range has no disposal can be implemented with a clock auction or a deferred acceptance heuristic.*

Proof. Given α, it can be implemented with a clock auction described as follows: Set $p(N) = \max B$ 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

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 auction. On the other hand, when the auction stops we have $A_t \subseteq \alpha\left(p_{A_t}\left(A^{t-1}\right), p_{N\setminus A_t}\left(A^{t-1}\right)^+\right)$, and putting together with the previous inclusion and using the no-disposal of the range of α implies $\alpha(b) = A_t = \alpha\left(p_{A_t}\left(A^{t-1}\right), p_{N\setminus A_t}\left(A^{t-1}\right)^+\right)$, hence the auction implements α. ■

Not all DA heuristics generate allocation rules that have substitutes. For example, the DA 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 $s_2^{\{1,2\}}(b_2) = s_1^{\{1\}}(b_2, b_1) = 0$ generates the allocation rue $\alpha(b_1, b_2) = \begin{cases} \{1,2\} & \text{if } b_1 < 1, \\ \varnothing & \text{otherwise,} \end{cases}$ which does not have subsitutes.

The following example shows that the no-disposal assumption is indispensable. Consider the allocation rule $\alpha(b) = \arg\min_{i\in N} b_i$, and suppose $B_1 = ... = B_N$ (so that ties exist, and in case of ties all of them are winners). Then there is no clock auction to implement α. For 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 to the reserve, then the reduction would eliminate any affected bidder. On the other hand, if some one bidder has bid least, then failing to reduce the price prevents the algorithm from ever identifying that bidder.

The no-disposal assumption is satisfied, e.g., in Example 1, if the feasible set F is *comprehensive* (meaning that $A \in F$ implies $A' \in F$ whenever $A \subseteq A'$) and heuristic has perfect feasibility checking. But not all heuristic rules satisfy it (e.g., in Example 5, the heuristic may accept a set $A' \subseteq N$ when bids are low and a set $A \subseteq A'$ when bids are high).

To see assumption of substitutes is not disposable, consider the efficient allocation rule in the example of Section 2, which is monotonic and whose range has no disposal, but which cannot be implemented by a clock auction.

12 Appendix C: Proof of Proposition 2

Let

$$
\begin{aligned}
F_{Zn} &= \{S \subseteq N : (\exists c : S \to \{1, ..., n\})(\forall s, s' \in S)(\{s, s'\} \in Z \Longrightarrow c(s) \neq c(s'))\}, \\
F^d_{mn} &= \{S \subseteq N : (\forall j)\, |S \cap N_j| \leq n - d\}, \text{ and } F^d_{Zmn} = F_{Zn} \cap F^d_{mn}.
\end{aligned}
$$

For each $S \subseteq N$, define $S_j = S \cap (\cup_{j' \leq j} N_{j'})$, $S'_j = S \cap N_j$, and $Z^*(S'_j) = \cap_{s \in S'_j}\{s' \in S_{j-1} : \{s, s'\} \in Z\}$ – i.e., the set of stations in S_{j-1} that interfere with *all* of the stations in S'_j. If $Z^*(S'_j) \subseteq N_{j-1}$, then to avoid interference it is necessary to assign a different channel to each station in $S'_j \cup Z^*(S'_j)$. A necessary condition for this is $\left|S'_j \cup Z^*(S'_j)\right| = \left|S'_j\right| + |Z^*(S'_j)| \leq n$, which is assumption (ii) of the proposition. Less obviously, Hall's Marriage theorem[24] implies that this condition is also sufficient, allowing us to prove the following lemma.

Lemma 3 $F^d_{Zmn} = F^d_{mn}$.

Proof. *It is obvious that $F^d_{Zmn} \subseteq F^d_{mn}$ (the first set imposes all the same within-area constraints plus additional ones). For the reverse inclusion, consider any $S \in F^d_{mn}$. We will establish that $S \in F^d_{Zmn}$ by showing the possibility of constructing the required channel assignment function $c : S \to \{1, ..., n\}$.*

Begin the construction by assigning a different channel $c(s)$ to each station $s \in S_1$, which is possible because $|S_1| \leq n$. Then, c is feasible for S_1. Inductively, suppose that the channel assignment c has been constructed to be feasible for stations S_{j-1}. We show how to extend c to a feasible channel assignment for $S_j = S_{j-1} \cup S'_j$.

[24] Hall's Marriage Theorem concerns bipartite graphs linking two sets – "men" and "women." Given any set of women S, let $A(S)$ be the set of men who linked ("acceptable") to at least one woman in S. Hall's theorem asserts that there exists a one-to-one match in which every woman is matched to some acceptable man (some men may be unmatched) if and only if for every subset S' of the women, $|S'| \leq |A(S')|$.

In our application, Hall's theorem is used to show that channels can be assigned to the stations in each area without violating the channel constraints implied by the assignments in the lower indexed areas.

For any $S''_j \subseteq S'_j$, define $J(S''_j) = \{1, ..., n\} - c(Z^*(S''_j))$ and notice that $|J(S''_j)| = n - |c(Z^*(S''_j))|$. According to Hall's Marriage Theorem (informally, think of the stations in S'_j as the women and the n channels as the men), there exists a one-to-one map $c : S'_j \to \{1, ..., n\}$ with the property that $(\forall s \in S'_j) c(s) \in J(\{s\})$ if and only if $(\forall S''_j \subseteq S'_j)$, $|J(S''_j)| \geq |S''_j|$. Substituting for $|J(S''_j)|$, this last inequality is equivalent to $|c(Z^*(S''_j))| + \left|S''_j\right| \leq n$.

For all $S''_j \subseteq S'_j$, since $S \in F^d_{mn}$, we have $\left|S''_j\right| \leq n - d$. Then, since $S''_j \cup S_{j-1} \in F^d_{mn}$, it follows from assumption (ii) of the Proposition that $|S''_j| + \left|Z^*(S''_j)\right| \leq n$. Combining that inequality with $\left|c(Z^*(S''_j))\right| \leq |Z^*(S''_j)|$, we obtain the condition required by Hall's Marriage theorem, implying the existence of a one-to-one function $c : S'_j \to \{1, ..., n\}$ such that $(\forall s \in S'_j) c(s) \in J(\{s\})$. This extends c to a feasible channel assignment for the expanded domain $S_j = S'_j \cup S_{j-1}$. ■

Finally, to establish the proposition, consider the following DA algorithm. At any round t, if the set of stations already assigned is T, then any station that is essential gets a score of zero. Among inessential stations at round t, the score for any station s with $m(s) = j$ is $n - |T \cap N_j| + b(s)/(1 + b(s))$. (Intuitively, this algorithm tries to keep the number of stations assigned in each area roughly equal at every round.) By the above Lemma, this algorithm will first assign the most valuable station in each area, then the second most valuable station in each area, and so on until at least the $n - d$ most valuable stations in each area are assigned.

13 Appendix D: Proof of Proposition 7

For the "if" direction of (i), recall from Proposition 3 that any assignment rule α that is implementable via a DA threshold auction is also implementable with a clock auction in which bidders use cutoff strategies with cutoffs corresponding to their sealed bids. Furthermore, we can compute the outcome of the *paid-as-bid* auction with assignment rule α using a "two-phase clock auction" mechanism, as follows. In phase 1, the clock auction algorithm 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," and then paying each winning bidder the last price it has accepted. In this two-phase clock auction game, if bidders use the cutoff profile b, that obviously leads to the same outcome as the paid-as-bid DA auction game with bid profile 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 payoff profiles surviving iterated deletion do not depend on the order of deletion. Hence, iterated deletion of dominated strategies in any order leads to the same set of possible outcomes. It will be convenient for us to also sometimes delete strategies that are *equivalent* to a surviving strategy, i.e., those that yield exactly the same auction outcomes (allocation and payments) as the surviving strategy for any remaining profile of others' strategies.[25] Obviously, deleting equivalent strategies does not affect the outcomes of iterated deletion, and so below we will use the term "b_i' $*$dominates b_i" to include both the case in which b_i' $*$dominates b_i and the case in which the two bids are equivalent.

We specify the following deletion process: Begin by deleting for each agent i all the bids/cutoffs $b_i < v_i^+$ (which are $*$dominated by 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 iteratively 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 $t = 1, 2, ...$ of our iterated deletion process, the set of strategies remaining to each bidder i is $\hat{B}_i^{t-1} = B_i \cap [v_i^+, \max\left\{v_i^+, p_i\left(A^{t-1}\right)\right\}]$. With just these strategies remaining, when the clock auction offers new prices

[25] Note that in a non-bossy paid-as-bid auction, two bids of bidder i are equivalent to each other against others' bids from $\hat{B}_{-i}$ if and only if they both make bidder i a sure loser, i.e., $i \notin \alpha\left(b_i, b_{-i}\right) \cup \alpha\left(b_i', b_{-i}\right)$ for all $b_{-i} \in \hat{B}_{-i}$.

$p(A^t)$ in iteration t, for each bidder i, all the cutoffs $b_i \in \hat{B}_i^{t-1}$ such that $b_i > \max\{v_i^+, p_i(A^t)\}$ are sure to lose and are therefore $*$dominated in the remaining subgame by bidding v_i^+. Deletion of these $*$dominated strategies yields new strategy sets $\hat{B}_i^t = B_i \cap [v_i^+, \max\{v_i^+, p_i(A^t)\}]$ for each i. These iterated deletions continue until phase 1 ends at some iteration T, at which the set of winners $\alpha(\hat{B}^T)$ is uniquely determined. For each agent i, if $\hat{B}_i^T$ is not a singleton, then its largest element, $\max \hat{B}_i^T = \max\{v_i^+, p_i(A^T)\}$, is $*$dominant in the remaining game with strategy sets $\hat{B}^T$ (because it wins at the highest remaining price). So, we may perform one more round of deletions, taking $\hat{B}_i^{T+1} = \{\max(v_i^+, p_i(A^T))\}$. Hence, the only possible outcome of iterative elimination of $*$dominated bids in any order is the outcome corresponding to the bid profile $(\max(v_i^+, p_i(A^T)))_{i\in N}$.

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 all bids of such bidder by the pure strategy bid v_i^+ to obtain another mixed strategy profile σ with the same distribution of outcomes.

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 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 $b = (\max\left\{v_i^+, p_i\left(A^T\right)\right\})_{i\in N}$ 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 $p_i\left(A^T\right) \geq v_i^+$, while each $i \in N\backslash A^T$ bids v_i^+, which is by definition above its threshold price. Thus by Proposition 6 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 paid-as-bid auction for allocation rule α is dominance solvable and construct a clock-auction price mapping $p : H \to \mathbb{R}^N$ that implements α. We construct p by iterating over possible clock auction histories and, in each iteration, reducing the price to a single bidder by a minimal decrement. For each possible clock auction history A^t of the auction and each bidder i, let $\hat{B}_i\left(A^t\right) \subseteq B_i$ denote the set of bidder i's bids (cutoffs) that are consistent with history A^t, i.e.,

$$\hat{B}_i\left(A^t\right) = \begin{cases} \{b_i \in B_i : b_i \leq p_i\left(A^{t-1}\right)\} & \text{for } i \in A_t, \\ \{(p_i\left(A^{t-1}\right))^+\} & \text{for } i \in N\backslash A_t. \end{cases}$$

We will show by induction that, for each possible history A^t, the strategy sets $\hat{B}_i\left(A^t\right)$ have two important properties: (a) $\cup_{b\in\hat{B}(A^t)}\alpha\left(b\right) \subseteq A_t$ (only bidders who are still active could become winners), and (b) the sets $\hat{B}\left(A^t\right)$ can be obtained by an iterative process that, at each, deletes some bids that are then always-losing bids for bidders in A_t or above some cutoffs for bidders in $N\backslash A_t$.

We initialize the clock prices at the null history N as $p(N) = \max B$, so that $\hat{B}_i(N) = B_i$ for each i and properties (a) and (b) are trivially satisfied. For each clock round $t = 1, 2, ...,$ given any history A^t at which properties (a) and (b) are satisfied, we show that either we can stop the auction at this point with the set of winners A_t, or we can reduce the price to a single bidder in such a way that properties (a) and (b) are inherited by any history A^{t+1} that could succeed A^t. We do so using the following claim:

Claim 1 *For all possible histories $A^t \in H$, either (i) $\alpha(b) = A_t$ for all $b \in \hat{B}(A^t)$ (all active bidders must win), or (ii) there exists $i \in A_t \setminus \cup_{b_{-i} \in \hat{B}_{-i}(A^t)} \alpha(p_i(A^{t-1}), b_{-i})$ (there is an active bidder whose highest remaining bid $p_i(A^{t-1})$ must lose).*

Proof. We begin by noting that if, given some history A^t, the set of winners is uniquely determined to be some $\hat{A} \subseteq N$ (i.e., $\hat{A} = \alpha(b)$ for all $b \in \hat{B}(A^t)$), then by inductive property (a), either $\hat{A} = A_t$ and so we are in case (i) of the claim, or we are in case (ii) of the claim for some bidder $i \in A_t \setminus \hat{A}$. Thus, it remains only to prove the claim for the case in which $\alpha(\hat{B}(A^t))$ is not a singleton.

Call two bids $b_i, b_i' \in \hat{B}_i(A^t)$ of bidder i *allocation-equivalent* (at A^t) if $\alpha(b_i', b_{-i}) = \alpha(b_i, b_{-i})$ for all $b_{-i} \in \hat{B}_{-i}(A^t)$. For each bidder i, we construct a strategy set $\bar{B}_i \subseteq \hat{B}_i(A^t)$ by eliminating from $\hat{B}_i(A^t)$ all of i's allocation-equivalent bids except for the highest one from each equivalence class. Note that by construction $\max \bar{B}_i = \max \hat{B}_i(A^t) = p_i(A^{t-1})$ for $i \in A_t$, and $\bar{B}_i = \hat{B}_i(A^t) = \left\{(p_i(A^{t-1}))^+\right\}$ for $i \in N \setminus A_t$. Note also that the elimination of allocation-equivalent bids preserves the possible sets of winners: $\alpha\left(\bar{B}\right) = \alpha(\hat{B}(A^t))$, and that, by assumption, this is not a singleton.

Now consider a value profile v such that $v_i^+ = \begin{cases} \min B_i & \text{for } i \in A_t, \\ (p_i(A^{t-1}))^+ & \text{for } i \in N \setminus A_t \end{cases}$ (thus, bidders in A_t always prefer to win, and the other bidders preferred to exit at the last prices they were offered). Observe that for this value profile, inductive property (b) allows us to obtain strategy sets $\hat{B}(A^t)$ by iterated deletion of strategies that are $*$dominated by bidding v_i^+, by deleting all bids below $(p_i(A^{t-1}))^+$ for all bidders $i \in N\backslash A_t$, and iteratively deleting always-losing bids. Next, note that when a bid b_i is allocation-equivalent to a bid $b_i' > b_i$, then b_i is $*$dominated by b_i'. Iterated deletion of such $*$dominated strategies from $\hat{B}(A^t)$ yields $\bar{B}$.

Dominance solvability for value profile v implies that if the set $\alpha(\bar{B})$ of winners is not uniquely determined, then for some bidder $i \in A_t$ some bid $b_i \in \bar{B}_i$ must be $*$dominated by some bid $b_i' \in \bar{B}_i \backslash \{b_i\}$ against $\bar{B}_{-i}$. If we had $b_i' > b_i$, then (by monotonicity) the two bids would have to win against the same set of opposing bid profiles $b_{-i} \in \bar{B}_{-i}$ and hence (by non-bossiness) they would be allocation-equivalent, which is impossible given our construction of $\bar{B}$.[26] Hence, we must have $b_i' < b_i$. Furthermore, since $\bar{B}$ was obtained from $\hat{B}(A^t)$ by deleting allocation-equivalent bids, bid b_i' must also $*$dominate bid b_i against $\hat{B}_{-i}(A^t)$. Since $v_i^+ = \min B_i \le b_i' < b_i$, such $*$dominance is only possible if b_i never wins against $\hat{B}_{-i}(A^t)$, which, by monotonicity, implies that the bid $p_i(A^{t-1}) = \max \bar{B}_i \ge b_i$ also never wins against $\hat{B}_{-i}(A^t)$. This establishes the claim. ■

Now, if we are in case (i) of the claim, then the auction can be finished in round t: if all the bidders have used cutoff strategies with a cutoff profile $b \in \hat{B}(A^t)$, then the auction has found the desired allocation $\alpha(b) = A_t$.

If we are instead in case (ii) of the claim, then in iteration t of the clock auction, our construction decrements the price to the bidder i identifed in the claim and leave the other prices unchanged, that is, we set $p_j(A^t) = \begin{cases} (p_i(A^{t-1}))^- & \text{for } i = j, \\ p_j(A^{t-1}) & \text{for } j \neq i. \end{cases}$. It remains to show that the two inductive properties are inherited in iteration $t+1$ by both the history $A^{t+1} = (A^t, A_t)$ in which bidder i accepts the reduced price and the history $A^{t+1} =$

[26] This argument relies on the observation that deleting allocation-equivalent bids for one player does not affect the allocation-equivalence of other players' bids: hence, when two bids of bidder i are allocation-equivalent against $\bar{B}_{-i}$, they must also be allocation-equivalent against $\hat{B}_{-i}(A_t)$.

$(A^t, A_t \backslash \{i\})$ in which bidder i quits. For property (a), using the fact that $\hat{B}(A^{t+1}) \subseteq \hat{B}(A^t)$ and the inductive hypothesis, we see that $\cup_{b\in\hat{B}(A^{t+1})}\alpha(b) \subseteq \cup_{b\in\hat{B}(A^t)}\alpha(b) \subseteq A_t$, which establishes the property for history $A^{t+1} = (A^t, A_t)$; as for history $A^{t+1} = (A^t, A_t \backslash \{i\})$, we use in addition the fact that $i \notin \cup_{b\in\hat{B}(A^{t+1})}\alpha(b)$ since we are in case (ii) of the claim. For property (b), it extends to history (A^t, A_t) since $\hat{B}_i(A^t, A_t) = \hat{B}_i(A^t) \backslash \{p_i(A^{t-1})\}$ and we are in case (ii) of the claim, and it extends to history $(A^t, A_t \backslash \{i\})$ since $\hat{B}_i(A^t, A_t \backslash \{i\}) = \left\{b_i \in \hat{B}_i(A^t) : b_i \geq p_i(A^{t-1})\right\}$.