

Clock Auctions and Radio Spectrum Reallocation *

Paul Milgrom Ilya Segal

Stanford University

*The authors thank Alexey Kushnir, Mohammad Akbarpour, Piotr Dworzak, Tim Roughgarden, Inbal Talgam Cohen, Simon Loertscher, Joshua Thurston-Milgrom, and Alex Wolitzky for their comments and suggestions. The authors are members of the Auctionomics' team that advised the US Federal Communications Commission concerning the design of the "Incentive Auction." Segal gratefully acknowledges the support of the Toulouse Network for Information Technology. Milgrom acknowledges support from the US National Science Foundation, grant number 1525730.

Abstract

We study the class of multi-round, multi-product clock procurement auctions that reduce offered prices at each round. When prices stop declining, the remaining bidders become the winning sellers. For single-minded bidders, each such auction has five properties not shared by Vickrey auctions: Each is obviously strategy-proof and group strategy-proof, sets prices that are Nash equilibrium winning bids in the related first-price auction, preserves winner privacy about values, and can be extended to satisfy a budget constraint. In simulations of the US Incentive Auction, a heuristic clock auction from this class achieves quick computations, high efficiency, and low prices.

1 Introduction

In April, 2017, the U.S. Federal Communications Commission concluded its so-called “Incentive Auction,” which reallocated the radio frequencies previously used for UHF TV channels 38-51 to wireless broadband services. The auction, authorized by a 2012 act of Congress,¹ involved purchasing broadcast rights from some TV stations in a competitive “reverse auction” at a cost of about \$10.1 billion, retuning (“repacking”) the remaining over-the-air broadcasters to operate in the remaining channels, and selling wireless broadband licenses in the cleared spectrum in a “forward auction” for about \$19.8 billion. This paper reports theoretical and simulation analyses related to the design of the reverse auction.^{2,3}

One characteristic that made FCC’s reverse auction design particularly challenging was the computational complexity of the underlying economic problem. The TV stations remaining on air after the auction (which turned out to be all but 175 of the 2,990 pre-auction U.S. and Canadian stations) needed to be repacked in a way that satisfies more than a million constraints. Each constraint precludes some pair of (geographically close) stations from broadcasting on the same or adjacent channels to avoid excessive interference between them, or limits the set of channels that a station may use. Just checking whether it is feasible to repack any *given* set of stations into any set of TV channels is an NP-hard problem, which means roughly that for any algorithm, the worst-case computation time grows exponentially with the problem size.⁴ FCC computational experiments showed that the problem of identifying the *maximum-value* feasible subset of stations could not be solved exactly in reasonable

¹Middle Class Tax Relief and Job Creation Act of 2012, Pub. L. No. 112-96, §§ 6402, 6403, 125 Stat. 156 (2012).

²For information about the forward auction and how the forward and reverse auctions interact, see Milgrom and Segal (2017).

³A related descending-clock auction design, which also uses a declining “base clock price” to set different prices for individual items in a series of rounds, has now been used by the FCC for its Connect America Fund “CAF II” auction which began in July 2018. That auction was not to buy spectrum, but to procure commitments to supply broadband services to underserved areas of the United States.

⁴This problem is referred to as the “frequency assignment” problem, and is a generalization of the “graph coloring problem.” See Aardal et al. (2007) for a survey of computational approaches to this problem. The exponential time claim depends on a widely believed but still unproven hypothesis in complexity theory, namely, that $P \neq NP$.

time.⁵

The intractability of exact computations poses a challenge for traditional auction designs. For example, it is tempting in this context to use a Vickrey auction, because of its well-known efficiency and incentive properties. However, the Vickrey auction proved deficient in several ways, not least of which is that the difficulty of optimization in this problem made computing even approximate Vickrey prices impossible. With 2000 stations, a 1% error in computing one of the two maximum values in the pricing equation would lead to a pricing error equal to 20 times the average station value, which is far too high for a practical auction mechanism.

To overcome the computational and other disadvantages of the Vickrey auction design, the U.S. Federal Communications Commission (2015) instead adopted a descending-clock auction design (proposed by Milgrom et al. 2012). This multi-round auction proposes a series of tentative prices, reducing some prices at each round and leaving others unchanged. Whenever a bidder’s price is reduced, that bidder may exit the auction and be repacked into the TV band. When the auction stops, the bidders that have not exited become winners and receive their last clock price. To ensure that the final set of losing bidders can be feasibly repacked into the available channels, the FCC auction reduced the price offer to a bidder only if it could identify a feasible assignment that adds the bidder to the set of stations to be repacked in the TV band. As described above, each feasibility checking step requires solving a large-scale NP-hard problem and the auction required about 75,000 such steps, so it was expected that the checker would sometimes fail to determine, in the allotted time, whether a set of stations can be feasibly repacked.⁶ When the checker fails to prove that a set of stations can be feasibly repacked, the auction treats it as if the repacking were infeasible.

Within this general description, one can vary the price-setting rule and the information

⁵The 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 computational constraints of economic allocation mechanisms, the formal economic literature, motivated by Hayek (1945) has focused on modeling communication costs (e.g., Hurwicz 1977, Mount and Reiter 1974, Segal 2007), which are trivial in the setting of single-minded bidders considered in this paper.

⁶Using recent advances in machine learning and certain problem-specific innovations, Frechette et al. (2015) developed a feasibility checker that solves more than 99% of the problems in auction simulations within 2 seconds, ensuring that time-outs would be rare.

revealed to bidders to design an auction that balances multiple objectives, including efficiency, cost minimization, computational feasibility, budget balance and others. This paper formulates the general class of these auctions and examines their theoretical properties and how a particular auction can be tailored for an application like the Incentive Auction. For concreteness, we focus on a procurement auction like the FCC’s reverse auction, which offers descending prices to sellers, but the same analysis also applies with obvious sign adjustments to selling auctions that offer ascending prices to buyers, as well as to double auctions that offer prices to both buyers and sellers.

We call these mechanisms “deferred-acceptance (DA) clock auctions” to highlight their connection to the famous deferred acceptance algorithm of Gale and Shapley (1982), in which rejections are irreversible but final acceptance is deferred to the end of the process.⁷ In a DA clock auction, a price reduction to a seller amounts to irreversible rejection of the offer to sell at the previous price, but permits the seller to remain active by agreeing to a new, lower price.⁸ For brevity, we will sometimes use the term “DA auction” to refer to a DA clock auction, although we will later introduce corresponding direct mechanisms to be called “direct DA auctions.”

Most previous studies of clock auctions have focused on settings in which the auctioneer’s optimization is computationally tractable, bidders are substitutes and all provisionally losing bidders are offered nearly the same price. The substitutes condition implies that the auctioneer would never regret rejecting an offer after other offers are improved, and therefore a DA clock auction that decrements prices slightly at each round to some or all provisionally losing bidders eventually results in an allocation that is efficient and stable. This paper

⁷As noted by Hatfield and Milgrom (2005), the Gale-Shapley algorithm modified to a setting with monetary transfers (as in Kelso and Crawford (1982) and Demange et al. (1986)) and applied to the case of a single buyer and multiple sellers can be viewed as a clock auction.

⁸We do not consider clock-based auctions that are not DA auctions. One such example is the “Dutch clock auction” – a descending-price selling auction, which is strategically equivalent to a pay-as-bid auction and is not strategy-proof. Another is Ausubel and Milgrom’s (2002) “cumulative-offer” clock auction, which sometimes “recalls” exited bids, and is not strategy-proof. A third example is the heuristic “clock auction” proposed by Lehmann et al. (2002), in which truthful bidding is a Nash equilibrium but not an obviously dominant strategy.

considers a wider set of clock auctions, generating a wider set of mappings from values to allocations (“allocation rules”), including ones in which bidders are sometimes complements, but our incentive analysis is narrower in limiting attention to single-minded bidders. As an example of complements, a DA auction can be designed so that the auctioneer will buy from either both bidders A and B or from neither.⁹ We show that any DA auction rule like that one, which does not satisfy the substitutes condition, cannot exactly maximize welfare or minimize expected costs. Nevertheless, even when DA clock auctions are not optimizing, they may still have a number of appealing properties.

One such property is a strengthening of traditional strategy-proofness: not only is it a dominant strategy for a bidder in the auction to bid truthfully (that is, to keep bidding while its price offer is at least its value and exit immediately afterwards), but truthful bidding is optimal even if the bidder does not understand the auctioneer’s price reduction rule or does not trust the auctioneer to adhere to the rule. To conclude that truthful bidding is optimal, all a bidder needs to know is that its clock price can never be increased, and that it can exit any time the price is reduced. One way to formalize this strong property is using Li’s (2017) notion of an “obviously dominant strategy,” which requires that for any alternative strategy, at any information set in the game for which the alternative strategy prescribes a different action, the best-case payoff from the alternative strategy against possible strategies of the other players must be no greater than the worst-case payoff from the obviously dominant strategy. In a DA auction, truthful bidding guarantees a nonnegative payoff starting from any information set, while any deviation from it involves either exiting at a price weakly above value or continuing at a price below value, so the payoff starting from that information set is nonpositive.¹⁰

⁹A clock auction can guarantee this by specifying that if either A or B exits, the other’s clock price falls to zero, ensuring that it will exit, too.

¹⁰In contrast, in a sealed-bid Vickrey auction, truthful bidding is not obviously dominant. The auction has just one information set, and the best payoff from any bid above value is strictly positive while the worst payoff from a truthful bid is zero. The failure of obviousness may explain why in laboratory experiments involving even the simplest single-object Vickrey auctions, bidders often fail to bid truthfully, while they do bid truthfully in the English auction (see Kagel et al. (1987)).

The obvious dominance of truthful bidding implies a second valuable property: that no *coalition* of bidders could deviate from truthful bidding in a way that makes all of its members strictly better off. For proof, note that obvious dominance implies that the first coalition member to deviate from truthful bidding cannot benefit strictly from the coalitional deviation. Thus, truthful bidding is a strong Nash equilibrium of a DA auction, and the social choice function implemented by the auction is weakly group strategy-proof.¹¹ This observation extends the results of Moulin (1999), Juarez (2009), Mehta et al. (2007), and Demange and Gale (1985), Demange et al. (1986), and Hatfield and Milgrom (2005) to a broader class of mechanisms.

A third property is that any DA auction can be modified to respect the auctioneer’s budget constraint by adding rounds in which prices continue to fall. In the Incentive Auction, the budget constraint was that the cost of buying broadcast rights could not exceed the forward auction revenue net of certain expenses and targets. To ensure that the budget was satisfied, the rules specified that if the constraint was not satisfied, then the number of channels to be cleared would be reduced and bidding would resume. If reducing the clearing target repeatedly still did not result in satisfying the budget constraint, then the auction would be cancelled.^{12,13} In contrast to the Vickrey auction, which cannot accommodate budget constraints, any DA auction can be extended in this way to create a budget-respecting extension. This extension, which is itself a DA clock auction with all the properties that implies, results in the same outcome as the unextended auction whenever that satisfies the budget constraint, but continues to ensure that the constraint is eventually satisfied. This property of respecting the auctioneer’s budget constraint with a strategy-proof mechanism generalizes the findings of several other papers.¹⁴

¹¹It is not *strongly* group strategy-proof: since winners’ price are determined by losing bids, thus a weakly Pareto-improving deviation could be achieved by a loser deviating to raise a winner’s price.

¹²In the actual Incentive Auction, the clearing target was reduced three times before the budget constraint was satisfied.

¹³Introducing a budget constraint causes the auctioneer’s substitutes condition to fail, because if an auction is cancelled when a seller A raises its bid, then A ’s price increase can reduce the demand for other sellers.

¹⁴Other known mechanisms that satisfy a budget constraint include cost-sharing mechanisms of Moulin (1999), Juarez (2009), and Mehta et al. (2007), double auctions of McAfee (1982), and procurement auc-

A fourth property of DA auctions with both practical and theoretical significance is that, in contrast to any direct mechanism, winning bidders in a DA clock auction reveal only the minimal information about their values needed to prove that they should be winning. We call this property “unconditional winner privacy (UWP).”¹⁵ The *practical* significance of UWP is that it may alleviate winners’ concerns about misuse of the revealed information, encourage participation by bidders who find it costly to figure out their exact values, and allow the auctioneer to conceal winners’ politically sensitive “windfalls.”¹⁶ The *theoretical* significance of UWP is conveyed by the following characterization theorem: a monotonic allocation rule can be computed by a communication protocol that preserves UWP if and only if it can be computed by a DA clock auction.

DA auction allocation rules can also be characterized by the algorithms to compute them. An allocation rule can be implemented by some DA clock auction with truthful bidding if and only if it can also be implemented by a “greedy rejection algorithm,” which iteratively rejects the least attractive remaining bids according to some bid-ranking criterion that may depend on the bid amount and information about the previously rejected bids and bidders.¹⁷

When can DA clock auctions produce nearly efficient outcomes, that is, ones that nearly minimize the total value of the acquired bidders? For our answer, we call upon well known results about the performance of greedy algorithms, and on simulations related to the Incentive Auction.

tions by Ensthaler and Giebe (2009, 2014). All of these mechanisms can be shown to be DA auctions. Building on an earlier version of this paper, Duetting et al. (2014b) have constructed approximately optimal budget-constrained DA double auctions, and Jarman and Meisner (2015) have shown that optimal budget-constrained procurement auctions can be implemented as DA auctions.

¹⁵It is a variation of the notion of “unconditional privacy” used in computer science (Brandt and Sandholm (2005)).

¹⁶For discussions of the importance of privacy in auctions, see Ausubel (2004) and McMillan (1994).

¹⁷These *greedy rejection heuristics are similar to the heuristic* previously proposed to create computationally feasible incentive-compatible mechanisms in the field of “Algorithmic Mechanism Design” (as pioneered by Lehmann et al. (2002), and also used by Mu’alem and Nisan (2008), Babaioff and Blumrosen (2008), and others), but with a crucial difference: the previously proposed heuristics greedily *accept* the most attractive bids as determined by some ranking criterion. Both greedy acceptance and greedy rejection algorithms, when paired with threshold pricing, lead to computationally simple strategy-proof auction mechanisms, but these two kinds of auctions have different strategic properties: auctions based on greedy acceptance do not have an obviously strategy-proof implementation and are not group strategy-proof. See Appendix A for a simple example illustrating these points.

- When the sets of bids that can be feasibly rejected are the independent sets of a matroid, a greedy rejection algorithm achieves the optimum. This outcome is replicated by a descending DA clock auction that offers the same price to all the bidders that can still be feasibly rejected.¹⁸
- When the only constraints are ones to limit the sum of the “sizes” of the rejected bidders, then the problem of optimally rejecting the highest cost bidders is a “knapsack problem,” for which optimization is NP-hard. However, the maximum can often be well approximated using the Dantzig greedy heuristic, and the same outcome can be implemented by a DA auction, as illustrated in the next section.
- Following an earlier draft of this paper, Duetting et al. (2014a) and Kim (2015) have explored other cases in which the DA auctions lead to “high” efficiency.
- In five simulations of a scaled-down version of the FCC’s problem by Newman et al. (2017), described in Section 9, we find that a DA auction with the pricing rule and feasibility checker used by the FCC achieved within 10% of the minimal value loss from repacking (average approximation was 5%), while incurring 14%-30% *lower* costs than the Vickrey auction (average cost savings was 24%), and using only a small fraction of the Vickrey auction’s computation time.¹⁹

By seeking to minimize the total virtual costs of acquired stations, instead of their total values, DA auctions can also sometimes achieve exact or approximate expected *cost* minimization. We show that if bidders’ values are independently drawn from known “regular” distributions and the feasible sets of rejected bids are the independent sets of a matroid, then there exists a DA auction that minimizes the expected total payment by the auctioneer.

If bidders’ values are instead independently drawn from an *unknown* distribution, then a DA clock auction can incorporate “yardstick competition” among bidders to reduce expected

¹⁸This finding is related to, but distinct from, the “forward” clock auction proposed by Bikhchandani et al. (2011) for *selling* the bases of a matroid, which uses a greedy acceptance algorithm.

¹⁹Even on the scaled-down problem, the average Vickrey computation took on average over 90 CPU days, while an FCC auction simulation took less than 1.5 CPU hours.

costs, as in Segal (2003). In follow-on work, Loertscher and Marx (2015) construct a sequence of DA clock auctions that achieves asymptotically optimal profits for a broker in a two-sided market with unknown distributions of values and costs, as the number of participants grows.

Another way to assess the cost performance of a DA auction is to compare its prices to those of a related *competitive equilibrium*, and to those of a related *paid-as-bid auction*. To evaluate problems with computational constraints that may make optimization intractable, we compute the auctioneer’s demand at any given vector of prices by applying the auction’s allocation rule. Then, a competitive equilibrium is an allocation and prices such that (i) each seller whose price strictly exceeds its cost sells its good, (ii) each seller whose price is strictly less than its cost does not sell, and (iii) the sellers from whom the auctioneer chooses to buy are the same as those that choose to sell. We show that the DA clock auction allocation coupled with the winners’ auction prices and prices for losing bidders equal to their values is a competitive equilibrium, and in fact a maximal-price equilibrium sustaining the auction’s allocation. These competitive equilibrium prices are also a full-information Nash equilibrium bid profile of the sealed, paid-as-bid auction that selects winners in the same way.^{20,21} In contrast, Vickrey prices are not generally the prices of any competitive equilibrium, and may be higher than the winning bids in any Nash equilibrium of the associated paid-as-bid auction.²²

The paper is organized as follows: Section 2 contrasts the performance of DA clock

²⁰A stronger version of this equivalence obtains when the auction rule prescribes the set of winners is unchanged when losers’ bids are increased. (This property, which holds when winners are selected by optimization, also holds for various other allocation rules.) We show in an Appendix that the full-information paid-as-bid auction using a DA allocation rule with this property is dominance-solvable: iterated deletion of weakly dominated strategies in any order yields a unique outcome. Furthermore, this dominance-solvability characterizes DA allocation rules.

²¹When bidders are substitutes, the Vickrey outcome can be implemented by a DA auction. Hence, our equivalence result can be viewed as extending the finding by Bernheim and Whinston (1986) that the coalition-proof equilibrium outcome of an optimizing paid-as-bid auction coincides with the Vickrey (and DA auction) outcome when bidders are substitutes.

²²For example, all the full-information Nash equilibria selected by the criteria of Bernheim and Whinston (1986) in the selling version of their problem have weakly lower prices than the Vickrey prices, and may have strictly lower prices when bidders are not substitutes. This low revenue/high cost problem of the Vickrey auction, observed by Ausubel and Milgrom (2006), motivated the use of “core-selecting auctions” (Day and Milgrom 2008), which sacrifice strategy-proofness. The present paper proposes a different solution to the problem, which preserves strategy-proofness but sacrifices outcome efficiency.

auctions and the Vickrey auction in a simple three-bidder example in which bidders may not be substitutes. Section 3 gives a formal definition of DA clock auctions and shows that, for any mechanism in this class, truthful bidding strategies are obviously dominant, and truthful bidding is a strong Nash equilibrium of the game. Section 4 characterizes the DA clock auctions as the set of auctions that preserve winners' unconditional privacy. Section 5 introduces direct-revelation deferred-acceptance auctions and shows that they are strategically equivalent to DA clock auctions. Section 6 provides examples of DA auctions that may be useful in practice. Section 7 shows that if some DA auction optimizes an objective resembling efficiency or cost, then it treats bidders as substitutes. Consequently, the treatment of complementary bidders is always a heuristic, non-optimizing one. Section 8 compares the cost performance of DA auctions to two theoretical standards: competitive equilibrium and full information equilibrium of paid-as-bid auctions. Section 9 describes simulations of the performance of the FCC's reverse auction rules. Section 10 discusses multi-minded bidders and their roles in the Incentive Auction. Section 11 concludes.

2 A Simple Example

Here, we illustrate properties of DA auctions with a simple example with three TV stations (labelled 1,2,3) and a single channel available in which to assign losing bidders. The three stations' values for their broadcast rights are denoted v_1, v_2, v_3 , respectively. We consider two possible cases, corresponding to different sets of constraints on the stations that can be assigned to continue broadcasting. An efficient assignment maximizes the total value of the stations that continue to broadcast, or equivalently minimizes the total value of the stations whose rights are purchased.

In the first case, no two stations can be assigned to the same channel, so the efficient outcome is for the most valuable station to be assigned and for the broadcast rights of the other two to be purchased. The Vickrey auction accomplishes that while paying the two less

valuable stations a price equal to the value of the most valuable station. A descending DA clock auction that offers the same price to all stations and reduces it until some station exits replicates the Vickrey outcome: the most valuable station exits at a price equal to its value, and the other two stations are then acquired at that price. More generally, this single-price DA auction replicates the Vickrey outcome when stations are substitutes or, equivalently, when the feasible sets of rejected bidders are the independent sets of a matroid (Milgrom, 2017).

In the second case, the three stations are arrayed in order along a line segment. The peripheral stations 1 and 3 can both broadcast on a channel without interfering with each other, but neither can share a channel with station 2. Thus, it is feasible to either acquire station 2 or stations 1 and 3 together. Notice that this example can be viewed as a special case of the knapsack problem, if a station’s “size” is defined to be the number of interference links it has (so that stations 1 and 3 each has size 1 and station 2 has size 2), and a collection of stations can be feasibly assigned to continue broadcasting if and only if the sum of their sizes does not exceed 2 (the “size of the knapsack”).²³

In this example, a simple DA auction similar to the FCC’s auction might operate as follows. The auction is guided by a single *base clock price* of q that starts high and declines continuously. Each station i is offered a price equal to $w_i q$, where the “weight” w_i is set by the auctioneer as some function of the station’s observable characteristics. Assuming that stations bid truthfully, then if $v_2/w_2 > \max(v_1/w_1, v_3/w_3)$, station 2 exits first and the auctioneer buys stations 1 and 3; while if the inequality is reversed, either station 1 or station 3 exits first, at which point station 2 can no longer be feasibly assigned, so its price is “frozen.” The third station is then induced to exit by running its price down to zero. In particular, if we let $w_1 = w_3 = 1$ and $w_2 = 2$, then the auction implements Dantzig’s (1957) greedy algorithm for the knapsack problem.

²³The “knapsack problem” is the problem of maximizing the sum of values of selected items subject to the single constraint that the sum of the “sizes” of the items may not exceed some constant (the size of the knapsack).

Clearly, there is no vector of weights (w_1, w_2, w_3) for which the DA auction would be guaranteed to yield an efficient outcome, which in this simple example is easy to compute: buy stations 1 and 3 if $v_1 + v_3 < v_2$ and buy station 2 if the inequality is reversed. While the Vickrey auction is efficient and the DA auction is not, the latter also has some advantages. First, any DA auction is obviously strategy-proof and weakly group strategy-proof. In contrast, the Vickrey auction is not group strategy-proof: for example, when it acquires stations 1 and 3, it pays them prices $v_2 - v_3$ and $v_2 - v_1$, respectively, and bidders 1 and 3 could cooperate to increase those by bidding less than their station values. Also, unlike the DA auction, the Vickrey price for each winner depends on the other winner's value, so implementing the Vickrey payment requires full revelation of the winners' values.

Depending on the values, DA auctions can also have a cost advantage over the Vickrey auction. For example, when the Vickrey auction acquires stations 1 and 3, it pays a total cost of $2v_2 - v_1 - v_3 > v_2$. This outcome is not in the core (it would be blocked by the coalition consisting of the buyer and bidder 2), and it is more costly than any full-information pure Nash equilibrium of the associated first-price auction.²⁴ It follows that the prices are inconsistent with any competitive equilibrium. In contrast, when the DA auction buys stations 1 and 3, its cost is $\frac{w_1+w_3}{w_2}v_2 \leq v_2$, provided that $w_1 + w_3 \leq w_2$, so this outcome is in the core. When bidder 2 wins the DA auction, its cost could sometimes be higher than that of Vickrey, but we will show that the price is nevertheless “competitive” in the sense that if the auctioneer used its (inefficient) allocation rule to select items at the last auction prices that bidders had accepted, it would make the same choices.

Furthermore, DA auctions can be geared towards minimizing the expected procurement cost in the style of Myerson (1981). For assume that bidder values are independently distributed with known distributions F_i , and for each bidder i , construct the virtual cost function $\gamma_i(v) = v + F_i(v)/F_i'(v)$ and assume that each is increasing. In the first example, in

²⁴In the general setting, Vickrey payments to bidders 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)).

which any one of the three stations can be assigned, the following DA auction is an “optimal auction.” There is a continuously declining base clock price q and the corresponding clock price for each station i is $\gamma_i^{-1}(q)$. If the auction begins with a high value of q , then the bidder with the highest virtual cost will be the first to exit, and so the mechanism will exactly minimize expected procurement cost. This construction has some generality: a similar auction minimizes expected costs when the stations that are feasible to repack are the independent sets of a matroid. In the second (“knapsack problem”) example, we could similarly set i ’s clock price to be $\gamma_i^{-1}(q/w_i)$ so long as i can feasibly exit. This auction is not optimal, but it replicates the performance of the Dantzig greedy algorithm in approximately minimizing the sum of the winning bidders’ virtual costs for each value realization, and therefore the sum of their expected payments.²⁵

Like every DA auction, this one can also be modified to satisfy a budget constraint or to cancel the auction if that is impossible. For example, suppose that the total procurement budget is B . If the total prices in some DA auction exceeds B , then the auction can simply be extended so that prices continue to fall until the total does not exceed B . If the offers at those prices are unacceptable, then prices can continue to decrease until the outcome is satisfactory or all prices are zero (and all stations exit). The budget-respecting extension is still a DA auction. In contrast, if the Vickrey auction must be cancelled when the total cost exceeds B , the resulting auction is not strategy-proof: if bidders 1 and 3 would be winners but for the budget constraint, each of them could profitably increase its own bid, thereby reducing the other’s Vickrey price and making cancellation less likely.

²⁵In the actual FCC Incentive Auction, when the “base clock price” was q , the price offered to any feasible UHF television station i was $p_i(q) = q(w_i \cdot Pop_i)^{1/2}$, where Pop_i is the station’s broadcast area (used as a proxy for its value distribution) and w_i is its interference count, with $w_i^{1/2}$ used as a proxy for the station’s “size.”

3 Clock Auctions and Truthful Bidding

We consider procurement mechanisms with N bidders, in which each bidder can either “win” (which means that his bid to supply a given good or bundle of goods is “accepted”) or “lose” (which means that his bid is rejected). We restrict attention to mechanisms in which winning bidders receive payments but losing bidders do not, which we henceforth refer to as “auctions.”

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 $v_i \in \mathbb{R}_+$ that makes him indifferent between winning and losing; we call v_i his “value.”²⁶

Informally, a Deferred Acceptance (“DA”) clock auction is a dynamic mechanism that, in a sequence of rounds, presents a nonincreasing sequence of prices to each bidder. Each bidder whose price is reduced in a round may choose to exit or continue. Bidders that have not exited are called “active.” Bidders who choose to continue when their prices are reduced are said to “accept” the lower price. When the auction ends, the remaining active bidders become the winners and are paid their last (lowest) accepted prices. Different DA auctions are distinguished by their pricing functions, which determine the sequence of prices presented.²⁷

To avoid technical complications, we restrict attention to auctions with discrete time periods, indexed by $t = 1, \dots$. The set of active bidders in period t is denoted by $A_t \subseteq N$. A period- t history consists of the sets of active bidders in all periods up to period t : $A^t = (A_1, \dots, A_t)$ such that $A_t \subseteq \dots \subseteq A_1$. Let H denote the set of all such histories for all possible $t \geq 1$. A descending DA clock auction is described by a price mapping $p : H \rightarrow \mathbb{R}^N$ such that $p(A^t) \leq p(A^{t-1})$ for all $t \geq 2$ and all A^t .

²⁶For 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.

²⁷The 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.

The DA auction gives rise to an extensive-form game among the registered participants, as follows. The bidders who register for the auction comprise the initially active set: $A_1 = N$. These bidders are committed to accept the opening prices, which are $p(N)$. In each period $t \geq 1$, given history A^t , the auction offers a profile of prices $p(A^t)$ to the bidders. If $t \geq 2$ and $p(A^t) = p(A^{t-1})$, the auction stops; bidder i is then a winner if and only if $i \in A_t$ and in that case i is paid $p_i(A^t)$. If $t \geq 2$ and $p_i(A^t) < p_i(A^{t-1})$, then i may either exit or accept the new price. Letting $E_t \subseteq A_t$ denote the set of bidders who exit, the auction continues in period $t+1$ with the new set of active bidders $A_{t+1} = A_t \setminus E_t$ and the new history $A^{t+1} = (A^t, A_{t+1})$. We restrict attention to auctions, in which the set $\{p(A^t)\}_{h \in H}$ of possible price offers is finite (ensuring that the auction ends in a bounded number of periods).

To complete the description of the extensive-form game, we need to describe bidders' information sets, given by functions $I_i : H \rightarrow S_i$. We allow arbitrary information sets except that we require that each bidder i observe his current price. Formally, for any two histories $h, h' \in H$, if $I_i(h) = I_i(h')$, then $p_i(h) = p_i(h')$.

The *truthful strategy* of agent i with value v_i in a DA auction accepts at history h if and only if $p_i(h) \geq v_i$. Given our assumptions, this strategy is measurable with respect to the agent's information. We can prove that the truthful strategy is not only a dominant strategy, but also an obviously dominant strategy in the sense of Li (2017):

Definition 1 (Li 2017) *Strategy σ_i of agent i is obviously dominant if, for any alternative strategy σ'_i , at any information set I_i of the agent at which σ'_i and σ_i prescribe different actions, the agent's payoff achieved by σ_i and any strategy profile σ_{-i} of other players such that (σ_i, σ_{-i}) visits I_i is at least as high as his payoff achieved by σ'_i and any strategy profile σ'_{-i} of other players such that $(\sigma'_i, \sigma'_{-i})$ also visits I_i .*

Proposition 1 *The truthful strategy is an obviously dominant strategy in a DA clock auction.*²⁸

²⁸This proposition corresponds to an informal statement in an early version of this paper, which has been formalized by Li (2017) using his definition of obviously dominant strategies.

Proof. A deviation from the truthful strategy involves either exiting at a price weakly above value or continuing at a price below value. Either deviation yields a nonpositive payoff for any behavior of others that is consistent with the information set, while the truthful strategy yields a nonnegative payoff for any behavior of others. ■

This proposition also implies that *coalitional* deviations from truthful bidding cannot be strictly Pareto improving:

Corollary 1 *In a DA auction, for every strategy profile σ , if all members of a coalition $S \subseteq N$ switch to truthful strategies then at least one member of S will receive a weakly higher payoff.*

Proof. Consider the first history $h \in H$ at which a player $i \in N$ bids non-truthfully under strategy profile σ . If all members of S switch to truthful strategies, then history h will also be reached, and agent i 's payoff will be weakly increased according to Proposition 1. ■

Thus, truthful strategies form a strong Nash equilibrium of any DA auction. Since truthful strategies do not condition on other bidders' values, it is an “ex post” strong Nash equilibrium, that is, a strong Bayesian-Nash equilibrium no matter what information bidders have about others' values.²⁹

4 Winners' Privacy

Here we formalize the notion that DA clock auctions are the only incentive-compatible communication protocols that preserve the privacy of winners. For this purpose, let $V_i \subseteq \mathbb{R}_+$ denote the set of bidder i 's possible values, and define a “communication protocol” to be an extensive-form game form, with each terminal node mapped to a set of auction winners, coupled with a mapping from each agent's value $v_i \in V_i$ to his strategy in the game form.

²⁹Like Vickrey auctions, DA clock auctions may also have Nash equilibria in which some bidders do not play their dominant strategies (for a description of the full set of equilibria of the Vickrey auction, see Blume, Heidhues, Lafky, Munster, and Zhang (2009)). However, we would not expect bidders to play weakly dominated strategies if they have sufficient uncertainty about each other.

The protocol implements an allocation rule that can be described by $\alpha : \prod_{i \in N} V_i \rightarrow 2^N$, where $\alpha(v) \subseteq N$ is the set of winning bidders in state $v \in V$.

Computer scientists (e.g., Brandt and Sandholm (2005)) say that a protocol satisfies “unconditional privacy” if no coalition of agents can infer any information about the other players’ values in the course of the protocol besides the information implicit in the final allocation.³⁰ We modify this definition in two ways. First, we weaken it by focusing only on *winners’* privacy. DA auctions do not preserve losers’ privacy, since their drop-out points can reveal a lot of information about their values. Second, we strengthen the definition to require that no information could be inferred about a winner’s value beyond that needed to establish that he himself should win (rather than establish the whole set of winners), even when the other $N - 1$ players pool their information:

Definition 2 *A communication protocol satisfies Unconditional Winner Privacy (UWP) if for any player i , any pair of his possible values $v_i, v'_i \in V_i$, and any values $v_{-i} \in V_{-i}$ of the other players such that the protocols results in bidder i winning in both states ($i \in \alpha(v_i, v_{-i}) \cap \alpha(v'_i, v_{-i})$), the protocol finishes in the same terminal node in states (v_i, v_{-i}) and (v'_i, v_{-i}) .*

Definition 3 *A communication protocol is Ex-Post Incentive Compatible (EPIC), if the prescribed strategies form a Nash equilibrium even if each bidder can observe the whole state $V = \prod_{i \in N} V_i$.*

Proposition 2 *A DA clock auction with truthful bidding satisfies UWP. Furthermore, any allocation rule on a finite state space V that is implementable with an EPIC communication protocol satisfying UWP is also implementable with a DA clock auction with truthful bidding.*

Proof. For the first statement, for player i to win in a DA clock auction with truthful bidding for two different values $v_i, v'_i \in V_i$ given any strategy profile of the other players,

³⁰This notion is also known as “non-cryptographic privacy,” since it permits neither private communication channels (as in private-key cryptography) nor agents’ computational constraints (as in public-key cryptography). A definition of privacy that allows such cryptographic tricks would not much restrict what can be implemented (see Izmalkov et al. (2005)).

player i must not exit in either case, and so the protocol must finish in the same terminal node in both cases.³¹

For the second statement, starting with an EPIC UWP protocol P implementing allocation rule α , we construct a DA auction with truthful bidding that also implements α . First, we note that since P is EPIC, the direct mechanism for α is dominant-strategy incentive compatible, and therefore α must be monotonic (see Section 5 below). Then, we construct the DA auction by induction on the histories of P , in such a way that the DA auction reveals at least as much information about players at each history as P does at its corresponding history. We initialize the opening price of each bidder i at $p_i(N) = \max V_i$ for each i (so that all types truthfully accept it). Then, at any history h of P , let $\bar{V} \subseteq V$ denote the set of states in which h is reached in P . By the usual “communication complexity” argument (e.g., Kushilevitz and Nisan (1997), Segal (2007)), this set must be a product set: $\bar{V} = \bar{V}_1 \times \dots \times \bar{V}_N$. Let i be the player who sends a message at history h in P . We replace his message with several rounds of the DA auction that reduce bidder i ’s price p_i to $p'_i = \max \{v_i \in \bar{V}_i \cup \{-1\} : v_i < p_i\}$ for as long as $i \notin \cup_{v_{-i} \in \bar{V}_{-i}} \alpha(p_i, v_{-i})$ – that is, as long as bidder i with value p_i could never win at history h in P .

At the end of the price reductions, all bidder i types who could not win at history h in P fully reveal themselves in the DA auction by exiting. If $p_i < \min \bar{V}_i$, then $i \notin \cup_{v \in \bar{V}} \alpha(v)$, that is, bidder i could never win at history h in P , and is sure to exit in the DA clock auction as his price is reduced to p_i . If instead $p_i \geq \min \bar{V}_i$, there exists $v_{-i} \in \bar{V}_{-i}$ such that $i \in \alpha(p_i, v_{-i})$, and by monotonicity of α we also have $i \in \alpha(v_i, v_{-i})$ for any type $v_i \leq p_i$, who has not exited the clock auction at this point, and therefore by UWP of P any such type must send the same message as type p_i at history h .

Thus, when bidder i responds truthfully to the price reductions, all bidder i types with values not exceeding the new price p_i do not reveal themselves in either the original protocol

³¹This argument actually establishes a somewhat stronger property than UWP: that even if the other players use non-truthful strategies (in computer science lingo, they are “Byzantine” rather than “selfish”), they still cannot learn additional information about player i ’s value without causing him to lose.

or the constructed DA auction, while all types with values above p_i fully reveal themselves in the DA auction by exiting. Hence the clock auction reveals at least as much information as P at the corresponding history. In each round of the clock auction, only bidders who could never win in P exit. At any terminal history h of P , $\alpha(\bar{V})$ must be a singleton, and so any bidders who could still win (that is, who have not exited the clock auction) must be winners. Thus, the constructed DA auction implements allocation rule α . ■

5 Equivalent Direct Mechanisms

Appealing to the Revelation Principle and the strategy-proofness of DA auctions, we can construct an equivalent direct revelation mechanism as follows. Let $V_i \subseteq \mathbb{R}_+$ denote the set of bidder i 's possible values. In the direct mechanism, each bidder reveals his value, and the mechanism implements an *allocation rule* $\alpha : \prod_{i \in N} V_i \rightarrow 2^N$ and a *payment rule* $\pi : \prod_{i \in N} V_i \rightarrow \mathbb{R}^N$ such that losing bidders are not paid, that is, $\pi_i(v) = 0$ for all $i \in N \setminus \alpha(v)$. Such triples $\langle V, \alpha, \pi \rangle$ may be called “direct auctions.”

Definition 4 *Allocation rule α is monotonic if $i \in \alpha(v_i, v_{-i})$ and $v'_i < v_i$ imply $i \in \alpha(v'_i, v_{-i})$.*

Definition 5 *A direct auction $\langle V, \alpha, \pi \rangle$ is a threshold auction if α is monotonic and the price paid to any winning bidder $i \in \alpha(v)$ is given by the threshold pricing formula:*

$$\pi_i(v_{-i}) = \sup \{v'_i \in V_i : i \in \alpha(v'_i, v_{-i})\}. \quad (1)$$

The following characterization of strategy-proof direct auctions is well known:

Proposition 3 *Any threshold auction is strategy-proof. Conversely, any strategy-proof direct auction has a monotonic allocation rule, and, if $V = \mathbb{R}_+^N$, it must be a threshold auction.*

We may describe the direct deferred-acceptance (DA) algorithm on value spaces V by a collection of *scoring functions* $(s_i^A)_{A \subseteq N, i \in A}$, where for each $A \subseteq N$ and each $i \in A$,

the function $s_i^A : V_i \times V_{N \setminus A} \rightarrow \mathbb{R}_+$ is nondecreasing in its first argument. The input to the algorithm is a value vector $v \in V$ and the algorithm processing works 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 $s_i^{A_t}(v_i, v_{N \setminus A_t}) = 0$ for all $i \in A_t$ then stop and output $\alpha(v) = A_t$; otherwise, let $A_{t+1} = A_t \setminus \arg \max_{i \in A_t} s_i^{A_t}(v_i, v_{N \setminus A_t})$ 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 value profile $v \in V$, when the algorithm stops the set of active bidders is exactly $\alpha(v)$.

By inspection, every DA algorithm 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 to the winners. This auction, like any threshold auction, is strategy-proof. Furthermore, the threshold prices can be computed in the course of the DA algorithm, by initializing the prices as $p_i^0 = +\infty$ for all i , and then updating them in each round $t \geq 1$ as follows:

$$p_i^t = \min \left\{ p_i^{t-1}, \sup \left\{ v'_i \in V_i : s_i^{A_t}(v'_i, v_{N \setminus A_t}) < s_j^{A_t}(v_j, v_{N \setminus A_t}) \text{ for } j \in A_t \setminus A_{t+1} \right\} \right\}$$

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 is the winner’s threshold price.

The next two results show that the direct mechanisms corresponding to clock auctions with truthful strategies are exactly direct DA threshold auctions.

Proposition 4 *The direct mechanism for a finite DA clock auction with state space $V = \mathbb{R}_+^N$ (more generally, such that $\{p_i(h) : h \in H\} \subseteq V_i \subseteq \mathbb{R}_+$) and truthful bidding is a DA threshold auction.*

Proof. Given a finite DA clock auction p , we construct the scoring rule for the direct DA auction in the following manner: Holding fixed a set of bidders $S \subseteq N$ and their values $v_S \subseteq \mathbb{R}^S$, let $A_t^S(v_S)$ denote the set of active bidders in round t of the clock auction if every

bidder from S bids truthfully and every bidder from $N \setminus S$ never exits.³² Now for any given $A \subseteq N$ and $i \in A$, define the score of agent i as the inverse of how long he would remain active in clock auction if he bids truthfully with value v_i and all bidders from $N \setminus A$ bid truthfully with values $v_{N \setminus A}$, while bidders in $A \setminus \{i\}$ never quit:

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

(Note that the score is $1/\infty = 0$ if agent i remains active for the remainder of the auction.) This score is by construction nondecreasing in v_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 ($\arg \max_{i \in A} s_i^A(v_i, v_{N \setminus A})$) is the set of bidders who would quit the soonest in the clock auction under truthful bidding. If no more bidders would ever exit the auction, then all active bidders have the score of zero, so the DA algorithm stops. Finally, each winner's final clock auction price is its threshold price: it would have lost by bidding any higher value in V_i in the DA auction, since this would correspond to rejecting the final price in the clock auction. ■

Proposition 5 *Every direct DA threshold auction with a finite state space V is a direct mechanism for some finite DA clock auction with truthful bidding.*

Proof. Given a direct DA threshold auction with a scoring rule s and a finite state space V , we construct an equivalent clock auction as follows. We begin with notation. Let $p_i^- = \max(V_i \cap (-\infty, p_i) \cup \{\min V_i - 1\})$ and $p_i^+ = \min(V_i \cap (p_i, +\infty) \cup \{\max V_i + 1\})$. These denote the result of decrementing or incrementing a price p_i by a minimal amount.

The auction then operates as follows. Set the opening prices to $p_i(N) = \max V_i$ for each i , so that all truthful bidders participate. In each subsequent round, set

³²Formally, initialize $A_1^S(v_S) = N$ and iterate by setting $A_{t+1}^S(v_S) = A_t^S(v_S) \setminus \{j \in S : v_j > p_j(A_t^S(v_S))\}$. The sequence $(A_t^S(v_S))_{t=1}^\infty$ must start repeating eventually (when the clock auction stops).

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

This decrements the price to every highest-scoring active bidder, where the scoring function is applied to the last accepted price. Because the auction maintains $p_i(A^t) = p_i(A^{t-1})$ for all $i \in N \setminus A_t$, it remembers the prices rejected by the bidders who exited, which are one decrement below their values.

Then equivalence is easy to see: First, for every history of the clock auction, under truthful bidding, the next set of bidders to exit consists of the bidders who have the maximum scores among the active bidders, 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 his highest value that would be still winning, that is, his threshold price. ■

These results imply that direct DA threshold auctions inherit strategy-proofness and weak group strategy-proofness from their clock auction counterparts (see Proposition 1 and Corollary 1, respectively).^{33,34} On the other hand, direct auctions lose *obvious* strategy-proofness and winners' privacy, making them less attractive for some applications.³⁵ At the same time, direct DA threshold auctions are a useful theoretical construct for understanding the objectives that may be achieved by means of clock auctions.

³³Note that such auctions are not generally *strongly* group strategy-proof, because a weak Pareto improving coalitional deviation may be obtained by a change in losing bids that increases a winner's threshold price.

³⁴The immediate implication is for DA threshold auctions on finite value spaces, but the arguments are easily extended to infinite value spaces.

³⁵A downside of clock auctions is that they are slower, often requiring many rounds to achieve good precision. It is also possible to specify "hybrid" designs that achieve a compromise between the speed of sealed-bid auctions and the attractive features of clock auctions. These hybrid designs include, for example, clock auctions with sealed intra-round bids (restricted to be between "start-of-round" and "end-of-round" prices), and the "survival auctions" of Fujishima et al. (1999).

6 Examples of Deferred-Acceptance Auctions

In this section we describe some examples that shed light on potential applications. These examples expand on the examples of section 2 as well as some observations made in the earlier literature and in follow-up work to an earlier version of this paper.

Example 1 (Ensuring Feasible Outcomes) *Suppose that it is only feasible for the auctioneer to accept bids of 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).³⁶ A DA clock auction maintains feasibility if it only reduces prices to stations i if $A \setminus \{i\} \in \mathcal{F}$, and only to one such station at a time.*

In the reverse auction of the FCC’s Incentive Auction, checking whether a given set $N \setminus A$ of the rejected bidders could be assigned to the available channels in a way that satisfies all interference constraints was an NP-hard problem which could not always be solved in the allotted time.³⁷ When the feasibility checker timed out without producing a “yes” or “no” answer, the time-out was treated as a “no” and the price offer for the station was not reduced. This guaranteed both feasibility and strategic simplicity for the bidders regardless of the precision of computations.

Example 2 (Optimization with Matroid Constraints) *Suppose that the goal is to find an “efficient” (social cost-minimizing) set of winning bids subject to a feasibility constraint, as follows:*

$$\alpha(v) \in \arg \min_{A \in \mathcal{F}} \sum_{i \in A} v_i. \quad (2)$$

Suppose further that we have a perfect feasibility checker, which can compute whether $A \in \mathcal{F}$. The simplest relevant scoring function is: $s_i^A(v_i, v_{N \setminus A}) = v_i \cdot 1_{A \setminus \{i\} \in \mathcal{F}}$. Set aside the issue of ties by assuming that the sets of possible values are disjoint: $V_i \cap V_j = \emptyset$ for $i \neq j$. Then, by a classical result in matroid theory (see Oxley (1992)), the resulting greedy rejection algorithm computes an efficient allocation rule α if and only if the feasible sets of rejected

³⁶E.g., in the example of Section 2, $\mathcal{F} = \{\{1, 2\}, \{3\}, \{1, 3\}, \{2, 3\}, \{1, 2, 3\}\}$.

³⁷In auction simulations, the customized feasibility-checking software developed for the incentive auction obtained the exact answer in 99.9% of the feasibility checking problems (Frechette et al. 2017).

bids ($\{R \subseteq N : N \setminus R \in \mathcal{F}\}$) are the independent sets of a matroid with ground set N . This algorithm is implemented by a clock auction that offers the same descending price to all the active bidders who could still be feasibly rejected, and “freezes” a bidder’s price when it can no longer be feasibly rejected.³⁸

The matroid property in Example 2 captures a notion of “one-for-one substitution” among bidders (in particular, it implies that all the maximal feasible sets of rejected bids – the “bases” of the matroid – have the same cardinality.) This substitution pattern 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. One way to accommodate that in the auction is illustrated in the following stylized example:

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

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

The problem of maximizing $\sum_{i \in N} v_i$ subject to $\sum_{i \in N} w_i \leq 1$ is known as the “knapsack problem.” Interpreting $w_i > 0$ as the “size” of bidder i , this problem is equivalent to problem (2) with the specified choice of \mathcal{F} .³⁹ This formulation generalizes the example of Section 2, so we again appeal to the famous “greedy” heuristic of Dantzig (1957), which, setting aside ties, corresponds to the direct DA algorithm with $s_i^A(v_i, v_{N \setminus A}) = v_i/w_i 1_{A \setminus i \in \mathcal{F}}$. This is equivalent

³⁸This is related to the analysis of Bikhchandani et al. (2011), who consider “selling” auctions in which the family of sets of bids that could be feasibly *accepted* is a matroid. Their proposed efficient clock auction increments prices for bidders who could still possibly lose, that is, ones whose rejection would preserve the spanning property of the set of active bidders. (This generalizes the clock auction proposed by Ausubel (2004).) Their auction implements the “greedy worst-out” heuristic algorithm, while our proposed reverse auction implements the “greedy best-in” heuristic algorithm. While either algorithm would yield efficient allocations in both the procurement and selling matroid auctions, only one of them constitutes a DA algorithm: in the procurement auction it is the best-in algorithm, and in the selling auction it is the worst-out algorithm.

³⁹In follow-work, 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.

to the clock auction in which a descending base clock price q determines the current price offer $w_i q$ to each active bidder i , and in which the price for bidder i “freezes” when the bidder can no longer be feasibly rejected. It can be shown that the inefficiency is no larger than the remaining empty space in the knapsack when the first item i is rejected multiplied by v_i/w_i .

For the FCC reverse-auction, one goal was to keep total procurement cost low: lower procurement costs contributed to more spectrum being cleared. That makes it interesting to generalize our section 2 example, showing how a DA auction can approximately minimize expected costs.

Example 4 (Expected Cost Minimization with Independent Values) *Suppose that bidders’ values v_i are independently drawn from distributions $F_i(v) = \Pr\{v_i \leq v\}$ on $V_i = [0, \bar{v}_i]$ for each i . 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 that 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 $\hat{s}_i^A(b_i, b_{N \setminus A}) = s_i^A \left(\gamma_i(b_i), (\gamma_j(b_j))_{j \in N \setminus A} \right)$.*

Example 5 (Expected Cost Minimization with Correlated Values) *Suppose that the auctioneer again cares about minimizing the expected total cost, but relax the assumption that bidders’ values are statistically independent. In such cases, as noted in Segal (2003), the auctioneer would optimally condition the price offered to one bidder on the bids of the others, including those that no longer have a chance of winning. Thus, in contrast to the preceding examples, in this context it can be helpful to condition scores on the values of already-rejected bids.*

For a simple example, if the auctioneer values acquiring each bidder at π and faces no feasibility constraints, it might use scoring functions $s_i^A(v_i, v_{N \setminus A}) = \max\{v_i - p_A^*(v_{N \setminus A}), 0\}$, where $p_A^*(v_{N \setminus A}) = \arg \max_p (\pi - p) \Pr\{v_i \leq p | v_{N \setminus A}\}$ is the optimal monopsony price for the posterior distribution of values given the rejected bids. In follow-up work to an earlier version of this paper, Loertscher and Marx (2015) show that the optimal expected profits can be approximated asymptotically with a DA clock auction for a large number of bidders whose values are drawn i.i.d. from an unknown distribution.

Another important objective in FCC’s “incentive auction” is satisfaction of a budget constraint: the reverse auction’s cost must be at least covered by the forward auction’s revenues. In general, rather than being fixed as in our example in Section 2, the available budget may depend on the set of items purchased.

Example 6 (Auctioneer’s Budget Constraint) *Suppose that the budget constraint is that when the set of winning bidders is A , the auctioneer is not permitted to pay more than $R(A)$ in total, with $R(\emptyset) = 0$. Any DA clock auction p can be modified to one that always satisfies the budget constraint and produces the same outcome as the original auction when the latter would satisfy the constraint, and otherwise may reduce the set of winning bidders or cancel the auction, if necessary. For this purpose, one simply changes p so that $p(A^t) < p(A^{t-1})$ whenever $\sum_{i \in A_t} p_i(A^{t-1}) > R(A_t)$.⁴⁰ The modified auction is still a DA auction, with all the properties that implies.*

In the mirror-image formulation of “selling” goods, budget-constrained “cost-sharing” DA clock auctions have been proposed by Moulin (1999) and Mehta et al. (2009). Their clock auctions always offer prices to all active bidders that exactly cover the cost of serving them (in our notation, $\sum_{i \in A_{t-1}} p_i(A^{t-1}) = R(A_{t-1})$ for all histories A^{t-1}), and stops as soon as all active bidders accept their prices (that is, $A_t = A_{t-1}$). A different kind of budget-constrained (sealed-bid and clock) DA auction for procurement is proposed by Ensthaler and Giebe (2009,

⁴⁰The equivalent direct DA threshold auction must use scoring that is contingent on already-rejected bids, since it is those bids that determine the current threshold prices of the still-active stations.

2014) for the case where the target revenue $R(A)$ does not depend on the set of winners A . In a follow-up to our work, Jarman and Meisner (2017) study optimal budget-constrained auctions for this case, and show that they can be implemented as DA auctions.

Budget constraints may be combined with various feasibility constraints on the set of accepted bids. For example, McAfee (1992) proposes a budget-constrained DA clock double auction for a homogeneous-good market with unit buyers and unit sellers, in which the feasibility constraint dictates that the number of the accepted buy bids (“demand”) must be equal to the number of the accepted sell bids (“supply”).⁴¹ 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. Nevertheless, the FCC’s setting also admits a double clock auction that, in the manner of McAfee’s double auction, sets a sequence of possible targets for the number of channels to clear, starting with the largest number and then reducing it whenever the revenue constraint cannot be satisfied with the current clearing target.⁴²

7 Optimization, Substitutes and Clock Auctions

A substantial literature on many-to-one matching (Gale and Shapley (1982), Kelso and Crawford (1982), Demange et al. (1986), Hatfield and Milgrom (2005)) has demonstrated that a stable allocation can be found using a deferred-acceptance algorithm in settings with an appropriate “substitutes” property, which guarantees that the responding agent will never regret rejecting a provisionally losing offer when other available offers have improved. (See

⁴¹McAfee’s clock auction offers the same ascending “buy” price to all buyers and the same descending “sell” price to all sellers, “freezing” the price to a “short” side of the market to keep demand within 1 of supply, and stopping as soon as both (i) the “sell” price falls weakly below the “buy” price and (ii) demand equals supply.

⁴²Duetting, 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.

Ausubel (2004) and Bikhchandani et al. (2007) for similar results for auctions.) In our setting, which is a special case of the ones considered in the literature (it has a single responding agent, and the proposing agents are single-minded), allocation rule α has *substitutes* if $i \in \alpha(v)$ and $v'_j > v_j$ for some $j \neq i$ implies $i \in \alpha(v'_j, v_{-j})$.

An allocation rule α that has substitutes on a finite state space can be implemented by a DA clock auction that decrements prices minimally only to those active bidders who wouldn't be accepted in α given current best offers (the active bidders' current prices and the exited bidders' last accepted prices), and continues until no such bidders can be found. Furthermore, in Appendix C we show that implementability with any clock auction satisfying those conditions on any finite product subspace characterizes the substitutes property.

While the substitute property is satisfied in some classical examples (such as Example 2 above), it is known to be quite restrictive (see, e.g., Hatfield and Milgrom (2005), Milgrom (2009), Ostrovsky and Leme (2015)). As discussed above, DA clock auctions can also achieve good performance for some practically important settings in which bidders cannot be treated as substitutes (such as settings with knapsack feasibility constraints, budget constraints, correlation in expected cost minimization). However, the proposed DA auctions for those settings were not exactly optimal, and we now show that this is not accidental: DA auctions are inconsistent with optimizing any objective function in a natural class. Thus, the handling of complements by DA clock auctions is always heuristic, not optimizing.

An allocation rule α is “optimizing” if

$$\alpha(v) \in \arg \max_{A \subseteq N} F(A) - \sum_{i \in A} \gamma_i(v_i) \quad (3)$$

for some function $F : 2^N \rightarrow \mathbb{R}^N \cup \{-\infty\}$ and functions $\gamma_i : \mathbb{R} \rightarrow \mathbb{R}$. For example, (3) is a surplus maximization problem when $\gamma_i(v_i) \equiv v_i$ and $F(A)$ is the auctioneer's gross benefit from accepting bid combination A (with infeasible combinations assigned $F(A) = -\infty$). Alternatively, (3) can describe maximization of the auctioneer's expected profit, or

minimization of its expected cost (as in Example 4), when bidders' values are independently drawn from regular distributions whose virtual values are given by $\gamma_i(v_i)$.

Proposition 6 *Suppose that allocation rule $\alpha : \mathbb{R}_+^N \rightarrow 2^N$ solves (3) for each $v \in \mathbb{R}_+^N$, for some nondecreasing continuous functions γ_i for each i . Further suppose that α restricted to any finite state space $V = \prod_{i \in N} V_i \subseteq \mathbb{R}_+^N$ is implementable with a DA clock auction. Then, α has substitutes on any such state space V on which the objective in (3) has no ties.*

The proof of the proposition is given in Appendix D.

8 Competitiveness of DA Auction Payments

In this section, we compare the prices in DA auctions with two notions of competitive prices, corresponding to competitive equilibrium and the full-information Nash equilibrium of the related paid-as-bid auction. For each of the concepts, the auctioneer uses the same DA allocation rule α to map bids to outcomes.

Definition 6 *Let $A \subseteq N$ and $p \in V$. We say that $\langle A, p \rangle$ is an α -competitive equilibrium in state $v \in V$ if (i) $v_i \leq p_i$ for all $i \in A$, (ii) $v_i \geq p_i$ for all $i \in N \setminus A$, and (iii) $A = \alpha(p)$.*

The following proposition establishes, for any DA allocation rule α , a three-way equivalence between (a) the winners' threshold prices $(\pi(v))_{i \in A}$ (which are also the DA auction prices), (b) the maximal competitive equilibrium prices supporting the winning allocation, and (c) winning bids in a full-information Nash equilibrium of the related first-price auction.

Proposition 7 *Let α be a DA allocation rule, and let $v \in V$, $A \subseteq N$ and $p_A \in V_A$. Then the following three statements are equivalent:*

- (a) $A = \alpha(v)$ and for each $i \in A$, $p_i = \pi_i(v_{-i})$,
- (b) $(p_A, v_{N \setminus A})$ is a maximal price vector such that $(A, p_A, v_{N \setminus A})$ is an α -competitive equilibrium in state v ,

(c) bid profile $(p_A, v_{N \setminus A})$ is a full-information Nash equilibrium of the paid-as-bid auction in state v , yielding allocation A .

Proof. (b) implies (a): We must have $A = \alpha(v) = \alpha(p_A, v_{N \setminus A})$ by equilibrium condition (iii) and UWP. Thus, by (1), $p_i \leq \pi_i(v_{-i})$ for all $i \in A$. This must hold with equality for the equilibrium to be maximal.

(a) and (b) imply (c): By (b), all bidders have nonnegative payoffs at bid profile $(p_A, v_{N \setminus A})$. It is trivial that bid reductions are not profitable for either a winner or a loser in the paid-as-bid auction. Since each winner is bidding its threshold price, any bid increase by a winner would cause him to lose. By monotonicity of α , a loser also cannot gain by increasing his bid.

(c) implies (b): Take any Nash equilibrium bid profile $p \in V$ with $p_{N \setminus A} = v_{N \setminus A}$ and let $A = \alpha(p)$. Then $\langle A, p \rangle$ satisfies competitive equilibrium conditions (ii) and (iii) by construction, and satisfies condition (i) due to winners' best response condition. Now, take any $p' \geq p$ such that $\langle A, p' \rangle$ is a competitive equilibrium in state v . Competitive equilibrium condition (ii) requires that $p'_{N \setminus A} \leq v_{N \setminus A} = p_{N \setminus A}$, and so $p'_{N \setminus A} = v_{N \setminus A}$. But we must also have $p'_A \leq p_A$, for otherwise a bidder $i \in A$ with $p'_i > p_i$ would have a profitable deviation to p'_i from bid profile p , remaining winning due to Unconditional Winners' Privacy. ■

As noted previously, there is no corresponding equivalence for the Vickrey auction: Vickrey prices may be neither competitive equilibrium prices nor consistent with any equilibrium of the corresponding paid-as-bid auction.

Proposition 7 leaves open the possibility that there may exist multiple competitive equilibrium outcomes, including ones in which losers' equilibrium prices are below their values. There could also be multiple Nash equilibrium outcomes of the first-price auction. In Appendix E we show that if the allocation rule α is non-bossy, these problems mostly vanish: there is a unique α -competitive allocation, which is also the unique allocation among undominated strategies in the related first-price auction game. In addition, the first-price auction game is then dominance solvable.

9 Simulations of the FCC Reverse Auction

The usefulness of the DA auctions introduced in this paper depends on their performance in complex applications. One way to examine that is to use simulations based on the FCC design and compare the outcomes to those of the VCG mechanism. Since fully optimizing the repacking of the 2,990 U.S. and Canadian stations to compute VCG allocations and prices has proven to be computationally intractable, in work with collaborators (Newman et al. 2017) we conducted auction simulations using a small enough set of stations for exact efficient channel assignments to be computed. The subset consists of 202 U.S. UHF stations starting with seven UHF stations in New York City and including all the UHF stations that are within two links from them in the interference graph, assuming that UHF stations are not interested in switching to VHF. This subset covers one of the densest areas of the U.S.: the average number of neighbors (stations with which a given station shares an interference constraint) in the subset is 46.1, and overall there are 70,384 channel-specific interference constraints (which are posted by FCC on site http://data.fcc.gov/download/incentive-auctions/Constraint_Files/). We simulated the problem of clearing 21 broadcast channels (126 MHz of spectrum), which was FCC’s initial clearing target in the Incentive Auction. Stations values for our simulations were drawn randomly using the methodology proposed by Doraszelski et al. (2016) based on publicly available data. The total value of the 202 stations in the dataset is \$13,309 billion.

We simulated both the VCG outcome and the outcome of the FCC auction, under the assumption that each station is owned separately and bids truthfully. For the FCC auction, we scored stations using the volumes used by FCC in the auction (posted in http://wireless.fcc.gov/auctions/incentive-auctions/Reverse_Auction_Opening_Prices_111215.xlsx).

We used the same computational environment (described in Newman et al. (2017)) for both approaches. On the scaled-down setting, the average VCG run took over 90 days of aggregate CPU time (the computation of VCG prices to the winners could be parallelized).

Then we simulated FCC’s auction for the same dataset, using the SATFC feasibility

checking software developed by Auctionomics, Inc. (described by Frechette et al. (2015)). We used SATFC version 2.3.1 (publicly released on <https://github.com/FCC/SATFC/releases>), which was also used by FCC in the actual auction. In contrast, the average FCC auction on the scaled-down setting took 1.5 hours of aggregate CPU time.⁴³

Since the VCG outcome minimizes the total value of stations taken off air, the corresponding value for the FCC auction design must necessarily be higher. We defined the *value loss ratio* of the DA mechanism as the sum of values of stations winning (and hence going off air) under the mechanism divided by the corresponding sum of values of winning stations in the efficient (VCG) allocation. In our simulations, the FCC auction design had efficiency loss ratios of between 0.8% and 10.0% (the average ratio was 5%). On the other hand, the costs (total payments to broadcasters) of the FCC auction were 14% to 30% *lower* than the costs paid by VCG (the average cost savings was 24%).

In addition, the FCC design proved to be computationally tractable and could be scaled up to the true size with roughly proportional increase in computation time: indeed, a nationwide computation could be run using about 6 CPU hours.⁴⁴ In contrast, sufficiently precise nationwide VCG computations were found to be impossible even in a matter of CPU weeks. Thus, the simulations established that, in contrast to the VCG auction, a heuristic DA auction with good feasibility checking is computationally feasible and can achieve significant cost savings by reducing stations' information rents, with smaller losses of allocative efficiency.

10 Multi-Minded Bidders

Our analysis has restricted bidders to be single-minded, that is, to bid for a single pre-determined option. The problem of algorithmic design of strategy-proof mechanisms for

⁴³With feasibility checking parallelized among 18 SATFC workers using 8 cores per worker, an actual simulation computing both the allocation and prices took about 7 minutes.

⁴⁴Even though nationwide simulations necessitated using time-outs on some feasibility checking problems, these were found to have no substantial effects on the overall efficiency and cost.

multi-minded bidders in computationally difficult environments is substantially more challenging. For example, in the DA algorithms of Gale and Shapley (1962) and Hatfield and Milgrom (2005), as well as their practical adaptations such as the Simultaneous Multi-Round Ascending Auctions used by the F.C.C. (Milgrom 2000, Gul and Stachetti (2000)), multi-unit buyers have incentives to engage in demand reduction (and similarly, multi-unit sellers have incentives to engage in supply reduction). While for some settings sophisticated DA clock auctions have been proposed in which truthful bidding is an ex post Nash equilibrium (Ausubel (2004), Ausubel et al. (2006), de Vries et al. (2007), Bikhchandani et al. (2011)), in those settings bidders are substitutes, computations are correspondingly easy, and the outcomes coincide exactly with the Vickrey outcome. For computationally challenging combinatorial-auction settings, proposed DA auctions include versions of sequential serial dictatorship (Bartal et al. 2003), randomized mechanisms that are “truthful in expectation” (Lavi and Swami 2011), and mechanisms that implement in undominated strategies (Babaioff et al. 2009). However, the approximation guarantees of these mechanisms shrink in the number of objects, so these mechanisms may not have worked well for the FCC reverse auction problem.

In the FCC reverse auction setting, a number of participants were potentially multi-minded: namely, they could have been interested in selling more than one station, or in switching to a VHF channel rather than giving up their licenses outright. Our recommendation to the FCC was based on our assessment that the crucial challenge of the auction was to attract participation of the hundreds of broadcasters with the lowest values, who would be the sellers in an efficient outcome. From the perspective of these mostly smaller bidders, this was an extremely complicated resource allocation, in which the rules for winner determination were barely comprehensible and certainly not computable. Few of these bidders would be likely to hire expensive consultants to advise their bidding. Most would plan to sell at most one station in any given market, and have only a single option in mind: selling their broadcast rights if the price was high enough. For such bidders, the obviously strategy-proof

DA clock auction would make it easy to decide to participate and bid.⁴⁵

11 Conclusion

The analysis and results reported in this paper were developed in connection with our engagement to advise the FCC on the design of its “Incentive Auction.” The search for pricing algorithms to lead to the best possible results inspired simulations that eventually led to the pricing rule described earlier. DA clock auctions have the huge advantage for this application that, in contrast to sealed-bid auctions, they are obviously strategy-proof and, in contrast to Vickrey auctions, are computable, weakly group strategy-proof, and compatible with auction budget constraints. Obvious strategy-proofness is important because it reduces the cost of participation, especially for small local broadcasters whose participation was believed to be critical for a successful Incentive Auction.⁴⁶

Our competitive pricing results highlight the good news that the obvious strategy-proofness that might attract participation by small broadcasters need not raise acquisition costs above the levels of competitive equilibrium prices or above the bids in a Nash equilibrium of the related paid-as-bid auction. This is important for efficiency as well, because the full Incentive Auction was a two-sided mechanism, in which the revenues from 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). High costs could lead to less clearing, less trade, and a loss of efficiency.⁴⁷

⁴⁵Given the DA clock auction design, market power by broadcasters was a potential problem which grew in importance after the auction design was substantially completed, as investors bought multiple television stations in individual markets. The significance of market power that may have resulted is an empirical question. By law, bids in the auction cannot be revealed and that limits our ability to assess the impact of supply withholding strategies in the actual reverse auction.

⁴⁶Two additional advantages make a DA clock auction particularly suitable for FCC’s problem and account for broad popularity of clock auctions in practice. First, clock auctions with information feedback can help bidders aggregate common-value information and thereby improve efficiency and revenues (as in Milgrom and Weber 1982). Second, clock auctions allow bidders to express their preferences for multiple bidding options (e.g., auction bundles) by switching among those options (expressing such preferences in a sealed bid may require a larger message space). However, modeling these advantages is beyond the scope of this paper.

⁴⁷We had also explored incorporating yardstick competition into the design, allowing the auction to set

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 a new class of auction designs and reaffirm the deferred-acceptance idea as the basis for some of the most successful new market mechanisms.

maximum prices for broadcasters in regions with little competition based on bids in other, more competitive regions. This option, however, was judged to be too complex.

12 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, which are 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 bidder 3’s value but strictly positive, they would both strictly benefit from both bidding above bidder 3’s value, 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 a DA clock auction, which is an ascending-price clock auction in the selling-auction setting, since the allocation is completely determined by which bidder has the highest value, while the first bidder exiting in a clock auction has the worst (lowest) value according to some scoring criterion.⁴⁸

⁴⁸Lehmann et al. (2002) propose a descending-price “clock” auction, in which, when bidder 3 buys first and the allocation is determined, the prices of bidders 1 and 2 must continue descending to determine the winner’s threshold payment. Note that in this “clock auction,” in contrast to the deferred-acceptance clock auctions studied in this paper, truthful bidding strategies are not dominant strategies (although they do form a Nash equilibrium).

13 Appendix B: A Near-Optimal DA Algorithm for TV Spectrum Repurchasing

Consider a setting in which 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 $\{1, \dots, n\}$ 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

$$\mathcal{F} = \{A \subseteq N : (\exists c : N \setminus A \rightarrow \{1, \dots, n\})(\forall i, j \in N \setminus A)(\{i, j\} \in Z \Rightarrow c(i) \neq c(j))\}. \quad (4)$$

Proposition 8 *Suppose there exists an ordered partition of the set N of stations into m disjoint sets N_1, \dots, N_m such that*

- (i) *for each $k = 1, \dots, m$ 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 $k = 1, \dots, m$ and each $S \subseteq N_k$ satisfying $|S| \leq n - d$, we have*

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

Consider a DA algorithm that iterates rejecting the most valuable bid in each partition element N_k as long as that is feasible and there are no more than $n - d$ stations rejected from each element, and then continues in any way. The total value of stations assigned by such an algorithm is at least a fraction $1 - d/n$ of the optimal value.

Proof. Let

$$F_{Zn} = \{S \subseteq N : (\exists c : S \rightarrow \{1, \dots, n\})(\forall s, s' \in S)(\{s, s'\} \in Z \implies c(s) \neq c(s'))\},$$

$$F_{mn}^d = \{S \subseteq N : (\forall j) |S \cap N_j| \leq n - d\}, \text{ and } F_{Zmn}^d = F_{Zn} \cap F_{mn}^d.$$

For each $S \subseteq N$, define $S_j = S \cap (\cup_{j' \leq j} N_{j'})$, $S'_j = S \cap N_j$, and $Z^*(S'_j) = \cap_{s \in S'_j} \{s' \in S_{j-1} : \{s, s'\} \in Z\}$ – that is, 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 $|S'_j \cup Z^*(S'_j)| = |S'_j| + |Z^*(S'_j)| \leq n$, which is assumption (ii) of the proposition. Less obviously, Hall’s Marriage theorem⁴⁹ implies that this condition is also sufficient, allowing us to prove the following lemma.

Lemma 1 $F_{Zmn}^d = F_{mn}^d$.

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

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 $S_j = S_{j-1} \cup S'_j$.

For any $S''_j \subseteq S'_j$, define $J(S''_j) = \{1, \dots, n\} - c(Z^*(S''_j))$ 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 $c : S''_j \rightarrow \{1, \dots, n\}$

⁴⁹Hall’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.

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 $|c(Z^*(S''_j))| + |S''_j| \leq n$.

For all $S''_j \subseteq S'_j$, since $S \in F_{mn}^d$, we have $|S''_j| \leq n - d$. Then, since $S''_j \cup S_{j-1} \in F_{mn}^d$, it follows from assumption (ii) of the Proposition that $|S''_j| + |Z^*(S''_j)| \leq n$. Combining that inequality with $|c(Z^*(S''_j))| \leq |Z^*(S''_j)|$, we obtain the condition required by Hall's Marriage theorem, implying the existence of a one-to-one function $c : S'_j \rightarrow \{1, \dots, n\}$ 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 $n - |T \cap N_j| + v(s)/(1 + v(s))$. (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. ■

Intuitively, the Proposition 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 .⁵⁰ 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

⁵⁰While 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.

(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 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.

14 Appendix C: Clock Auctions for Substitutes

We say that a clock auction is *consistent with* allocation rule α on a finite state space V if it initializes prices as $p_i(N) = \max V_i$ for each i , and then for every history A^t for $t > 1$, prices satisfy $p_i(A^t) \in \{p_i(A^t), p_i(A^{t-1})^-\}$, with $p_i(A^t) = p_i(A^{t-1})^-$ only for $i \in A_t \setminus \alpha(p_{A_t}(A^{t-1}), p_{N \setminus A_t}(A^{t-1})^+)$, and $p(A^t) = p(A^{t-1})$ only if $A_t \subseteq \alpha(p_{A_t}(A^{t-1}), p_{N \setminus A_t}(A^{t-1})^+)$. In words, the auction (i) decrements prices (minimally) only to those active bidders who wouldn't be accepted in α given current best offers (the active bidders' current prices and the exited bidders' last accepted prices), and (ii) continues until no such bidders can be found. Note that (i) ensures that a member of $\alpha(v)$ is not rejected when $(p_{A_t}(A^{t-1}), p_{N \setminus A_t}(A^{t-1})^+) = v$ (that is, all bidders' current best offers are their exact values), while (ii) ensures that a member of $N \setminus \alpha(v)$ is not accepted when the auction stops in such a situation.

Proposition 9 *Consider an allocation rule $\alpha : \mathbb{R}_+^N \rightarrow 2^N$. Allocation rule α restricted to any finite state space $V = \prod_{i \in N} V_i \subseteq \mathbb{R}_+^N$ is implemented by any clock auction consistent with*

α on V if and only if α is monotonic, has non-bossy winners, and has substitutes.⁵¹

Proof. The “if” part: First observe that an agent $i \in \alpha(v)$ could never exit the clock auction under truthful bidding. Indeed, at any history A^t at which $i \in A_t$ and i faces price $p_i(A^{t-1}) = v_i$, while for the other (truthful) bidders $p_{A_t \setminus \{i\}}(A^{t-1}) \geq v_{A_t \setminus \{i\}}$ and $p_{N \setminus A_t}(A^{t-1})^+ = v_{N \setminus A_t}$, the substitute property and $i \in \alpha(v)$ imply that $i \in \alpha(p_{A_t}(A^{t-1}), p_{N \setminus A_t}(A^{t-1})^+)$, and so the auction must set $p_i(A^t) = p_i(A^{t-1})$, ensuring that $i \in A_{t+1}$. Thus, we have $\alpha(v) \subseteq A_t$ throughout the auction under truthful bidding. In particular, this implies that $\alpha(p_{A_t}(A^{t-1}), p_{N \setminus A_t}(A^{t-1})^+) \subseteq A_t$, and so when the auction stops in period T we must have $\alpha(p_{A_T}(A^{T-1}), p_{N \setminus A_T}(A^{T-1})^+) = A_T$. Since under truthful bidding we have $p_{N \setminus A_T}(A^{T-1})^+ = v_{N \setminus A_T}$ and $v_{A_T} \leq p_{A_T}(A^{T-1})$, iteratively applying monotonicity and non-bossiness of winners for members of A_T yields $\alpha(v) = A_T$.

The “only if” part: Clearly, monotonicity and non-bossiness of winners are necessary for α to be implementable by *any* clock auction. To see the necessity of the substitutes condition, take any $v \in \mathbb{R}_+^N$, any $j \in N$, and any v'_j such that $v'_j > v_j$. Consider value spaces $V_j = \{v_j, v'_j\}$ and $V_k = \{v_k\}$ for all $k \neq j$, and the clock auction that starts with prices $p(N) = (v'_j, v_{-j})$ and then decrements the prices of all provisional losers (that is, members of $N \setminus \alpha(p(N))$). In order to implement α in state v it cannot decrement the price to any bidder from $\alpha(v) \setminus \{j\}$ in round 1, thus we must have $\alpha(v) \setminus \{j\} \subseteq \alpha(p(N)) \setminus \{j\} = \alpha(v'_j, v_{-j}) \setminus \{j\}$. ■

The assumptions of monotonicity and non-bossiness of winners are not dispensable: that is, they are not implied by substitutes.⁵² On the other hand, both assumptions are satisfied by any DA-implementable allocation rule.⁵³ Thus, an allocation rule that is implementable by *some* clock auction is implementable by *any* clock auction consistent with it if and only

⁵¹If we did not require the myopic clock auction to work on any value subspace (effectively, for any opening prices), then, as shown by Hatfield and Kojima (2009), weaker notions of substitutability may suffice.

⁵²For an allocation rule that has substitutes but bossy winners, let $N = 2$, $V_1 = V_2 = \{1, 2\}$, and $\alpha(v) = \arg \min_{i \in \{1, 2\}} v_i$ on value spaces $V_1 = V_2 = \{1, 2\}$ (so that in case of a tie for the lowest value both tied bidders are winners). Then α satisfies substitutes, but does not satisfy non-bossiness of winners: $1 \in \alpha(1, 2) \cap \alpha(2, 2)$ but $\alpha(1, 2) = \{1\} \neq \{1, 2\} = \alpha(2, 2)$.

⁵³In particular, non-bossiness of winners follows from winners' privacy, discussed in Section 4.

if the allocation rule has the substitute property.

15 Appendix D: Proof of Proposition 6

Before proving the proposition, we prove a useful lemma. Say that allocation rule α is *non-bossy* if for any $i \in N$, $v \in V$, and $v'_i \in V_i$, $\alpha(v'_i, v_{-i}) \cap \{i\} = \alpha(v) \cap \{i\}$ implies $\alpha(v'_i, v_{-i}) = \alpha(v)$ (that is, a bidder cannot affect the allocation without changing his own winning status.)

Lemma 2 *If allocation rule $\alpha : V \rightarrow 2^N$ is the unique solution to (3) for each $v \in V$, for some nondecreasing functions $\gamma_i : V_i \rightarrow \mathbb{R}$ for each i , then α is monotonic and non-bossy.*

Proof. For monotonicity, note that increasing bidder i 's value from v_i to $v'_i > v_i$ does not increase the objective (3) on any $A \subseteq N$ such that $i \in A$ and does not change the objective (3) on any $A \subseteq N \setminus \{i\}$.

For non-bossiness, note that (3) implies that for all $i \in N$, $v_i, v'_i \in V_i$, $v_{-i} \in V_{-i}$,

$$i \notin \alpha(v_i, v_{-i}) \cup \alpha(v'_i, v_{-i}) \implies \{\alpha(v_i, v_{-i})\} = \arg \max_{A \subseteq N: i \notin A} F(A) - \sum_{j \in A} \gamma_j(v_j) = \{\alpha(v'_i, v_{-i})\},$$

$$i \in \alpha(v_i, v_{-i}) \cap \alpha(v'_i, v_{-i}) \implies \{\alpha(v_i, v_{-i})\} = \arg \max_{A \subseteq N: i \in A} F(A) - \sum_{j \in A \setminus \{i\}} \gamma_j(v_j) = \{\alpha(v'_i, v_{-i})\}.$$

■

Proof of Proposition 6. Let $\Lambda(v) \subseteq 2^N$ denote the set of maximizers in (3) at value profile $v \in \mathbb{R}_+^N$. Suppose in negation that α violates substitutes when there are no ties: for two agents $i \neq j$, for some $v_i, v'_i \in \mathbb{R}_+$ such that $v_i < v'_i$, and some $v_{-i} \in \mathbb{R}_+^{N-1}$ we have $\Lambda(v_i, v_{-i}) = \{A\}$, $\Lambda(v'_i, v_{-i}) = \{A'\}$, with $j \in A \setminus A'$. Due to non-bossiness and monotonicity established in Lemma 2, this is only possible if we also have $i \in A \setminus A'$. By monotonicity and continuity of γ_i , there exists $\hat{v}_i \in (v_i, v'_i)$ such that $\Lambda(\hat{v}_i, v_{-i}) = \{A, A'\}$. Then, by continuity of γ_i, γ_j there exists $\varepsilon > 0$ small enough so that $\Lambda(\hat{v}_i + \delta_i, v_j + \delta_j, v_{-\{i,j\}}) \subseteq \{A, A'\}$ whenever $|\delta_i|, |\delta_j| \leq \varepsilon$. By monotonicity, we have $\Lambda(\hat{v}_i - \varepsilon, v_j - \varepsilon, v_{-\{i,j\}}) = \{A\}$, and using also

continuity of γ_i, γ_j , there exists $\delta \in (0, \varepsilon)$ small enough such that $\Lambda(\hat{v}_i + \delta, v_j - \varepsilon, v_{-\{i,j\}}) = \Lambda(\hat{v}_i - \varepsilon, v_j + \delta, v_{-\{i,j\}}) = \{A\}$. At the same time, by monotonicity, we have $\Lambda(\hat{v}_i + \delta, v_j + \delta, v_{-\{i,j\}}) = \{A'\}$. But then α cannot be implemented by any clock auction on value spaces

$$V_k = \begin{cases} \{\hat{v}_k - \varepsilon, \hat{v}_k + \delta\} & \text{for } k \in \{i, j\}, \\ \{v_k\} & \text{otherwise} \end{cases} \quad (\text{where we let } \hat{v}_j = v_j). \text{ Indeed, the clock auction}$$

cannot stop while $(p_i(h), p_j(h)) \geq (\hat{v}_i + \delta, \hat{v}_j + \delta)$, since bidders i and j must both lose when their values are $(\hat{v}_i + \delta, \hat{v}_j + \delta)$, but reducing the price to a bidder $k \in \{i, j\}$ below $\hat{v}_k + \delta$ would prevent him from winning in state $(\hat{v}_k + \delta, \hat{v}_{\{i,j\} \setminus k} - \varepsilon, v_{-\{i,j\}})$. ■

16 Appendix E: Competitive Equilibria and Paid-as-Bid Auctions with Non-Bossy Allocation Rules

In this Appendix, we sharpen the relationship between competitive equilibria and equilibrium of paid-as-bid auctions when the allocation rule is non-bossy. For one class of non-bossy rules, recall optimizing allocation rules to which Lemma 2 in Appendix D applies. For another, rules computed by DA algorithm with fixed scoring and perfect feasibility checking are non-bossy:

Lemma 3 *Suppose that $\mathcal{F} \subseteq 2^N$ is a comprehensive family of feasible sets (see Example 1), and that scoring is given by $s_i^A(v_i, v_{N \setminus A}) = \begin{cases} \sigma_i(v_i) & \text{if } A \setminus \{i\} \in \mathcal{F}, \\ 0 & \text{otherwise,} \end{cases}$ where $\sigma_i : V_i \rightarrow \mathbb{R}_{++}$ 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. Every DA procedure satisfies non-bossiness of winners (this is a corollary of Unconditional Winner Privacy established in Section 4). To check that condition for the losers, too, we show that if given value profile v agent i is rejected in some round t , then replacing his value with some $v'_i > v_i$ does not affect the final outcome of the algorithm. Note that it suffices to check situations in which the replacement results in bidder i is rejected *prior to*

round t : Indeed, otherwise he will be rejected in round t and the replacement will not affect the behavior of the algorithm. Suppose first that the replacement results in bidder 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 bidder j be the bidder rejected in round $t - 1$, we must have $A_{t-1} \setminus \{i, j\} \in \mathcal{F}$, and

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

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

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

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

First, we show that for non-bossy rules, there is a (unique) *maximum-price* competitive-equilibrium $\langle A, p \rangle$ in state v (that is, such that any other competitive equilibrium $\langle A', p' \rangle$ in that state has $p' \leq p$):

Proposition 10 *Let α be a non-bossy DA allocation rule. Then, letting $A = \alpha(v)$, $p_i = \pi_i(v_{-i})$ for each $i \in A$, and $p_{N \setminus A} = v_{N \setminus A}$, $\langle A, p \rangle$ is the maximum-price competitive equilibrium in state $v \in V$.*

Proof. By Proposition 7, $\langle A, p \rangle$ is a competitive equilibrium in state v . Now, let $\langle A', p' \rangle$ be another competitive equilibrium in state v . Using equilibrium conditions, monotonicity of α , non-bossiness, UWP, and (1), we have $A' = \alpha(p'_A, p'_{N \setminus A}) = \alpha(p'_A, v_{N \setminus A}) = \alpha(\pi_i(v_{-i}), v_{N \setminus A}) = A$. Then, equilibrium conditions and (1) imply $p'_i \leq \pi_i(v_{-i}) = p_i$ for all $i \in A$ and $p'_i \leq v_i$ for all $i \in NA$. ■

As for paid-as-bid auctions with non-bossy DA allocation rules, uniqueness of Nash equilibrium outcome can be ensured upon ruling out bids below values. To ensure that such bids are dominated strategies, we consider DA allocation rules with finite bid spaces V , and value profiles $v \in \mathbb{R}_+^N$ that are “generic,” defined as $v_i \in \mathbb{R}_+ \setminus V_i$ and $v_i < \max V_i$ for each i . Then, any bidder i ’s bid below his “rounded-up value” $v_i^+ = \min \{b_i \in V_i : b_i > v_i\}$ is either dominated by bidding v_i^+ or never wins.

Upon ruling out such dominated strategies, we obtain a single Nash equilibrium outcome. Furthermore, we can single out the same outcome by iterated deletion of dominated strategies, and we can show that this dominant-solvability in fact characterizes non-bossy DA rules.

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.*

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 11 *Consider a paid-as-bid auction with a monotonic, non-bossy allocation rule α and finite bid spaces.*

(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 $b_i = \max \{v_i^+, \pi_i(v_{-i}^+)\}$ for each $i \in N$, resulting in the equilibrium allocation $\alpha(b) = \alpha(v^+)$ and payments $p_i = \pi_i(v_{-i}^+)$ for all winners $i \in \alpha(b)$.

The equivalence of dominance solvability and implementation using a DA algorithm – equivalently, using the results of Section 5, a clock-auction algorithm – is constructive and intuitive. Indeed, note that in a round of iterated deletion of dominated strategies in a paid-as-bid auction, a bid is dominated by a *lower* bid if the former bid (and therefore any higher bid) would lose against all remaining possible bid profiles of the other bidders (while the latter might still win and is still above the bidder’s value). Deletion of bids that are dominated in this way corresponds to a clock auction’s decrementing of the price to a bidder who could no longer win at that price, given the prices already accepted by the other bidders. As for when a bid is dominated by a *higher* bid, this happens if (i) the former bid is below the bidder’s value (and so, this allows us to delete bids below value), or (ii) the two bids have the same chances of winning against all remaining opponents’ bid profiles (so, in particular, this allows us to delete sure winners’ bids below their threshold prices). Non-bossiness allows us to appeal to the results by Marx and Swinkels (1997) to choose a convenient order of iterated deletion without affecting the final outcome.

16.1 Proof of Proposition 11

For the “if” direction of (i), recall from Proposition 5 that any assignment rule α that is implementable via a DA threshold auction is also implementable with a clock auction. 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 V_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 each bidder i is restricted to use a truthful strategy corresponding to some value profile $b_i \in V_i$ (that is, exit when his price falls below b_i but not before), 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 the game satisfies the TDI condition of Marx and Swinkels (1997). Hence, by their results, the payoff profiles surviving iterated deletion of weakly dominated strategies do not depend on the order of deletion, and on whether we delete “very weakly” dominated strategies (which include payoff-equivalent strategies to a surviving strategy, in addition to those weakly dominated by it).⁵⁴ Hence, iterated deletion of either weakly or very weakly dominated strategies in any order leads to the same set of possible outcomes.

We specify the following deletion process of very weakly dominated (henceforth, VW-dominated) strategies: Begin by deleting for each agent i all the bids $b_i < v_i^+$ (which are VW-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 bid v_i^+ . At the beginning of each step $t = 1, 2, \dots$ of our iterated deletion process, the set of strategies remaining to each bidder i is $\hat{B}_i^{t-1} = B_i \cap [v_i^+, \max\{v_i^+, p_i(A^{t-1})\}]$. With just these strategies remaining, when the clock auction offers new prices $p(A^t)$ in iteration t , for each bidder i , all the bids $b_i \in \hat{B}_i^{t-1}$ such that $b_i > \max\{v_i^+, p_i(A^t)\}$ are sure to lose and are therefore VW-dominated in the remaining subgame by bidding v_i^+ . Deletion of these VW-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,

⁵⁴Note that in a non-bossy paid-as-bid auction, two bids of bidder i are payoff-equivalent to each other against others' bids from \hat{B}_{-i} if and only if they both make bidder i a sure loser, that is, $i \notin \alpha(b_i, b_{-i}) \cup \alpha(b'_i, b_{-i})$ for all $b_{-i} \in \hat{B}_{-i}$.

$\max \hat{B}_i^T = \max \{v_i^+, p_i(A^T)\}$, 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 VW-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 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 $b = (\max \{v_i^+, p_i(A^T)\})_{i \in N}$ 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 $p_i(A^T) \geq v_i^+$, while each $i \in N \setminus A^T$ bids v_i^+ , which is by definition above its threshold price. Thus by Proposition 7 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 $p : H \rightarrow \mathbb{R}^N$ that implements α . For each possible clock auction history A^t of the auction and each bidder i , let $\hat{B}_i(A^t) \subseteq B_i$ denote the set of bidder i 's bids that are consistent with history A^t , that is,

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

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

To construct the clock auction p , 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 $t = 1, 2, \dots$, 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 known to be 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 $\hat{A} \subseteq N$ (that is, $\hat{A} = \alpha(b)$ 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 $\bar{B}_i = \hat{B}_i(A^t) = \{(p_i(A^{t-1}))^+\}$ for $i \in N \setminus A_t$. Note also that the elimination of allocation-equivalent bids preserves the possible sets of winners: $\alpha(\bar{B}) = \alpha(\hat{B}(A^t))$, and that, by assumption, this is not a singleton.

Now consider a generic 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 VW-dominated by bidding v_i^+ , by deleting all bids below $(p_i(A^{t-1}))^+$ for all bidders $i \in N \setminus A_t$, 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 VW-dominated by b'_i . Iterated deletion of such VW-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 VW-dominated by some bid $b'_i \in \bar{B}_i \setminus \{b_i\}$ 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} .⁵⁵ 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 VW-dominate bid b_i against $\hat{B}_{-i}(A^t)$. Since $v_i^+ = \min B_i \leq b'_i < b_i$, such VW-dominance is only possible if b_i never wins against $\hat{B}_{-i}(A^t)$, which, by monotonicity, implies that the bid $p_i(A^{t-1}) = \max \bar{B}_i \geq b_i$ 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 the bid profile is $b \in \hat{B}(A^t)$, 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 identified 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 $A^{t+1} = (A^t, A_t \setminus \{i\})$ 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 $A^{t+1} = (A^t, A_t \setminus \{i\})$, 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 $\hat{B}_i(A^t, A_t) = \hat{B}_i(A^t) \setminus \{p_i(A^{t-1})\}$ and we are in case (ii) of the claim, and it extends to history $(A^t, A_t \setminus \{i\})$ since $\hat{B}_i(A^t, A_t \setminus \{i\}) = \{b_i \in \hat{B}_i(A^t) : b_i \geq p_i(A^{t-1})\}$.

⁵⁵This 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 $\hat{B}_{-i}(A_t)$.

References

- 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.
- Ausubel, L.M. (2004) “An efficient ascending-bid auction for multiple objects,” *American Economic Review*, 94 (5), 1452–1475.
- Ausubel, L.M. and P. Milgrom (2002), “Ascending Auctions with Package Bidding, ” *Frontiers of Theoretical Economics*, 1 (1), 1452–1475.
- 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.
- Ausubel, L., P. Cramton, and P. Milgrom (2006), “The Clock-Proxy Auction: A Practical Combinatorial Auction Design,” in *Combinatorial Auctions*, Peter Cramton, Yoav Shoham and Richard Steinberg (eds), MIT Press.
- Aygun, O., and T. Sonmez (2013), “Matching with Contracts: Comment,” *American Economic Review* 103(5), 2050-20451.
- Babaioff, M., and L. Blumrosen (2008), “Computationally-feasible truthful auctions for convex bundles,” *Games and Economic Behavior* 63(2), 588–620.
- Babaioff, M., R. Lavi, and E. Pavlov (2009), “Single-Value Combinatorial Auctions and Algorithmic Implementation in Undominated Strategies,” *Journal of the ACM* Volume 56(1).
- 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.

- Bernheim, D., and M. Whinston (1986), "Menu Auctions, Resource Allocation, and Economic Influence," *Quarterly Journal of Economics* 101(1): 1-31.
- 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.
- Blume, A., Heidhues, P., Lafky, J., Munster, J., Zhang, M. (2009), "All Equilibria of the Multiunit Vickrey Auction," *Games and Economic Behavior* 66, pp. 729–741.
- Brandt, F., and T. Sandholm (2005), "Unconditional Privacy in Social Choice" *Proceedings of Theoretical Aspects of Rationality and Knowledge (TARK)*.
- Bulow, J., J. Levin, and P. Milgrom (2009), "Winning Play in Spectrum Auctions," NBER Working Paper 14765.
- Dantzig, G. (1957), "Discrete-variable extremum problems," *Operations Research* 5, 266–277.
- Day, R. W., and P. Milgrom (2008), "Core-Selecting Package Auctions." *International Journal of Game Theory*, 36, 2008, 393–40.
- Demange, G., and D. Gale (1985), "The Strategy Structure of Two-Sided Matching Markets," *Econometrica* 53(4), pp. 873-888.
- Demange, G., D. Gale, and M. Sotomayor (1986), "Multi-Item Auctions," *Journal of Political Economy* 94(4), pp. 863-872.
- Dobzinski, S., N. Nisan, and Michael Schapira (2012), "Truthful randomized mechanisms for combinatorial auctions," *Journal of Computer and System Sciences* Volume 78 (1), Pages 15–25.

- Doraszelski, U., K. Seim, M. Sinkinson, and P. Wang (2016), “Ownership Concentration and Strategic Supply Reduction,” Centre for Economic Policy Research working paper DP11173.
- Duetting, P., V. Gkatzelis, and T. Roughgarden (2014a), “The Performance of Deferred-Acceptance Auctions,” *ACM Conference on Electronic Commerce*.
- Duetting, P., T. Roughgarden, and I. Talgam-Cohen (2014b), “Modularity and Greed in Double Auctions,” *ACM Conference on Electronic Commerce (ACM’ EC)*.
- 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.
- Ensthaler, L. and T. Giebe (2009): “Subsidies, Knapsack Auctions and Dantzig’s Greedy Heuristic,” SFB/TR 15 Discussion Paper No. 254.
- Ensthaler, L., and T. Giebe (2014), “A dynamic auction for multi-object procurement under a hard budget constraint,” *Research Policy* 43, 179-189.
- Federal Communications Commission (2015), “Public Notice: Broadcast Incentive Auction Scheduled to Begin on March 29, 2016.” FCC 15-78, AU Docket No. 14-252, GN Docket No. 12-268, WT Docket No. 12-269, MB Docket No. 15-146. Released August 11, 2015.
- Frechette, A., N. Newman, and K. Leyton-Brown (2015), “Deep Optimization for Spectrum Repacking,” *Communications of the ACM (CACM)*, to appear, 2017.
- Fujishima, Y., D. McAdams, and Y. Shoham (1999), “Speeding up ascending-bid auctions,” *Proceedings of the International Journal Conference in Artificial Intelligence*, 554-559.
- Gale, D., and L.S. Shapley (1962), “College Admissions and the Stability of Marriage,” *American Mathematical Monthly* 69, 9–14.

- Gul, F., and E. Stacchetti (2000), “The English auction with differentiated commodities,” *Journal of Economic Theory* 92(1), 66–95.
- Hayek, F.A. (1945): “The Use of Knowledge in Society,” *American Economic Review* 35, 519-30.
- Hatfield, J. and F. Kojima (2009), “Group Incentive Compatibility for Matching with Contracts,” *Games and Economic Behavior* 67: 745-749.
- Hatfield, J., and P. Milgrom (2005), “Matching with Contracts,” *American Economic Review* 95(4), 913-935.
- 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.
- Izmalkov, S., S. Micali, and M. Lepinski (2005), “Rational secure computation and ideal mechanism design,” *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 585-594.
- Jarman, F., and V. Meisner (2017), “Ex-post Optimal Knapsack Procurement,” *Journal of Economic Theory*, 171, 35-63.
- Juarez, R. (2009), “Prior-free cost sharing design: group strategy-proofness and the worst absolute loss,” in *Social Computing and Behavioral Modeling*, M. Young, J. Salerno and H. Liu (eds), Springer.
- 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.
- Kelso, A., and V. Crawford (1982), “Job Matching, Coalition Formation, and Gross Substitutes,” *Econometrica* 50(6), 1483-1504.

- 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.
- Kojima, F., P. Pathak and A. Roth (2013), “Matching with Couples: Stability and Incentives in Large Markets,” *Quarterly Journal of Economics*.
- Kim, A. (2015), “Welfare Maximization with Deferred Acceptance Auctions in Reallocation Problems,” 12rd European Symposium on Algorithms.
- Kushilevitz, E., and N. Nisan (1997), *Communication Complexity*, Cambridge University Press, Cambridge.
- Lavi, R., and C. Swami (2011), “Truthful and Near-Optimal Mechanism Design via Linear Programming,” *Journal of the ACM* 58 (6).
- Lehmann, D., L.I. O’Callaghan, Y. Shoham (2002), “Truth Revelation in Approximately Efficient Combinatorial Auctions,” *Journal of the ACM* 49(5), 577-602.
- Li, Shengwu (2017), “Obviously Strategy-Proof Mechanisms,” *American Economic Review*, 107(11): 3257-3287.
- List, J. and D. Lucking-Reiley (2000), “Demand Reduction in Multi-Unit Auctions: Evidence from a Sportscard Field Experiment,” *American Economic Review* 90(4), pp.961-972.
- Loertscher, S., and L.M. Marx (2015), “Prior-Free, Privacy Preserving Bayesian Optimal Mechanisms,” working paper.
- Marx, L.M., and J.M. Swinkels (1997), “Order Independence for Iterated Weak Dominance,” *Games and Economic Behavior* 18(2), 219–245.
- McAfee, R.P., (1992), “A Dominant Strategy Double Auction,” *Journal of Economic Theory* 56(2), 434–450.

- McMillan, J. (1994), "Selling Spectrum Rights," *Journal of Economics Perspectives*, 8(3), 145-162.
- Mehta, A., T. Roughgarden, and M. Sundararajan (2009), "Beyond Moulin Mechanisms," *Games and Economic Behavior* 67(1), 125–155.
- Milgrom, P. (2000), "Putting Auction Theory to Work: The Simultaneous Ascending Auction," *Journal of Political Economy* 108(2), pp. 245-272.
- Milgrom, P. (2004), *Putting Auction Theory to Work*, Cambridge University Press.
- Milgrom, P. (2009), "Assignment Messages and Exchanges," *American Economic Journal: Microeconomics*, 1(2), 95-113.
- Milgrom, P. (2017), *Discovering Prices: Auction Design in Markets with Complex Constraints*, Columbia University Press.
- Milgrom, P., L. Ausubel, J. Levin, and I. Segal (2012), "Incentive Auction Rules Option and Discussion," Appendix C to the FCC's Notice of Proposed Rulemaking, GN Docket No 12-268, October 2, 2012.
- Milgrom, P., I. Segal (2017), "Designing the US Incentive Auction (2017)," chapter 37 in *Handbook of Spectrum Auction Design*, Martin Bichler and Jacob Goeree (eds), Cambridge: Cambridge University Press.
- Mount, K., and S. Reiter (1974): "The Information Size of Message Spaces," *Journal of Economic Theory* 28, 1-18.
- Mu'alem, A., and N. Nisan (2008), "Truthful approximation mechanisms for restricted combinatorial auctions," *Games and Economic Behavior* 64(2), 612–631.
- Moulin, H. (1999), "Incremental cost sharing: Characterization by coalition strategy-proofness," *Social Choice and Welfare* 16, 279-320.

- Myerson, R. (1981), “Optimal Auction Design,” *Mathematics of Operations Research* 6, 58-73.
- Nisan, N., and A., Ronen (1999), “Algorithmic Mechanism Design,” *Proceedings of the thirty-first annual ACM Symposium on Theory of Computing (STOC)*, 129-140.
- Newman, N., K. Leyton-Brown, P. Milgrom, and I. Segal (2017), “Assessing Economic Outcomes in Simulated Reverse Clock Auctions for Radio Spectrum,” arXiv:1706.04324.
- Ostrovsky, M., and Leme (2015), “Gross Substitutes and Endowed Assignment Valuations,” working paper.
- Oxley, James G. (1992). *Matroid Theory*. Oxford University Press, Oxford, GB.
- Papadimitriou, C. and Y. Singer (2010): “Budget Feasible Mechanisms,” Working Paper (University of California at Berkeley).
- 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.
- Segal, I. (2003), “An Optimal Pricing Mechanism with Unknown Demand,” *American Economic Review* 93(3), 509-529.
- Segal, I. (2007), “The Communication Requirements of Social Choice Rules and Supporting Budget Sets,” *Journal of Economic Theory* 136, 341-378.
- Singer, Y. (2010), “Budget Feasible Mechanisms,” *51st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 765-774.
- 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.