Clock Auctions and Radio Spectrum Reallocation

Paul Milgrom

Stanford University

Ilya Segal

Stanford University

We study the class of multiround, multiproduct 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.

I. Introduction

In April 2017, the US Federal Communications Commission (FCC) concluded its so-called incentive auction, which reallocated the radio frequencies previously used for ultrahigh frequency (UHF) television

We thank Alexey Kushnir, Mohammad Akbarpour, Piotr Dworczak, 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 1525730).

Electronically published December 9, 2019
[*Journal of Political Economy*, 2020, vol. 128, no. 1]

channels 38–51 to wireless broadband services. The auction, authorized by a 2012 act of Congress, involved purchasing broadcast rights from some television 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.[1] This paper reports theoretical and simulation analyses related to the design of the reverse auction.[2]

One characteristic that made the FCC's reverse auction design particularly challenging was the computational complexity of the underlying economic problem. The television stations remaining on air after the auction (which turned out to be all but 175 of the 2,990 preauction US 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 television channels is an NP-hard problem, which means roughly that for any algorithm, the worst-case computation time grows exponentially with the problem size.[3] FCC computational experiments showed that the problem of identifying the maximum value feasible subset of stations could not be solved exactly in reasonable time.[4]

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 the least of which is that the difficulty of optimization in this problem made computing even approximate Vickrey prices impossible. With

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

[2] For information about the forward auction and how the forward and reverse auctions interact, see Milgrom and Segal (2017). A 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 Phase II auction, which ran in July–August 2018. That auction was not to buy spectrum but to procure commitments to supply broadband services to underserved areas of the United States.

[3] This 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 $P \neq NP$.

[4] The problem of designing computationally feasible economic mechanisms is studied in the field of algorithmic mechanism design (Nisan and Ronen 1999). While economists have long been concerned about computational constraints of economic allocation mechanisms, the formal economic literature (motivated by Hayek 1945) has focused on modeling communication costs (e.g., Mount and Reiter 1974; Hurwicz 1977; Segal 2007), which are trivial in the setting of single-minded bidders considered in this paper.

2,000 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 FCC (2015) instead adopted a descending-clock auction design (proposed by Milgrom et al. 2012). This multiround 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 television band. When the auction stops, the bidders who have not exited become winners and receive their last clock price. To ensure that the final set of losing bidders can be feasibly packed 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 television 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.[5] 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 revealed to bidders to design an auction that balances multiple objectives, including efficiency, cost minimization, computational feasibility, budget balance, and so on. 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 DA algorithm of Gale and Shapley (1962), in which rejections are irreversible but final acceptance is deferred to the end of the process.[6] In a DA clock auction, a price reduction

[5] Using recent advances in machine learning and certain problem-specific innovations, Frechette, Newman, and Leyton-Brown (2018) 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.

[6] As 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, Gale, and Sotomayor 1986) and applied to the case of a single buyer and multiple sellers can be viewed as a clock auction.

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.[7] 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 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.[8] 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 (i.e., to keep bidding while its price offer is at least its value and exit immediately afterward), 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 the bidder 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

[7] We 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, O'Callaghan, and Shoham (2002), in which truthful bidding is a Nash equilibrium but not an obviously dominant strategy.

[8] A 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.

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.[9]

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.[10] This observation extends the results of Demange and Gale (1985), Demange, Gale, and Sotomayor (1986), Moulin (1999), Hatfield and Milgrom (2005), Juarez (2009), and Mehta, Roughgarden, and Sundararajan (2009) 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.[11] 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 and otherwise continues to ensure that the constraint is eventually satisfied. This

[9] In 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, Harstad, and Levin 1987).

[10] It is not strongly group-strategy-proof, since winners' prices are determined by losing bids; thus, a weakly Pareto-improving deviation could be achieved by a loser deviating to raise a winner's price.

[11] In the actual incentive auction, the clearing target was reduced three times before the budget constraint was satisfied. Introducing a budget constraint causes the auctioneer's substitutes condition to fail because if an auction is cancelled when seller A raises its bid, then A's price increase can reduce the demand for other sellers.

property of respecting the auctioneer's budget constraint with a strategy-proof mechanism generalizes the findings of several other papers.[12]

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)*.[13] 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.[14] 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.

The 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.[15]

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 on well-known results about the performance of greedy algorithms and on simulations related to the incentive auction.

[12] Other known mechanisms that satisfy a budget constraint include cost-sharing mechanisms of Moulin (1999), Juarez (2009), and Mehta, Roughgarden, and Sundararajan (2009); double auctions of McAfee (1992); and procurement auctions 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, Roughgarden, and Talgam-Cohen (2014b) have constructed approximately optimal budget-constrained DA double auctions, and Jarman and Meisner (2017) have shown that optimal budget-constrained procurement auctions can be implemented as DA auctions.

[13] It is a variation of the notion of unconditional privacy used in computer science (Brandt and Sandholm 2005).

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

[15] These 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, O'Callaghan, and Shoham 2002 and also used by Babaioff and Blumrosen 2008, Mu'alem and Nisan 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 app. A (apps. A–E are available online) for a simple example illustrating these points.

1. 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 who can still be feasibly rejected.[16]
2. 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 (1957) greedy heuristic, and the same outcome can be implemented by a DA auction, as illustrated in the next section.
3. Following an earlier draft of this paper, Duetting, Gkatzelis, and Roughgarden (2014a) and Kim (2015) have explored other cases in which the DA auctions lead to high efficiency.
4. In five simulations of a scaled-down version of the FCC's problem by Newman et al. (2017), described in section IX, 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.[17]

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 costs, as in the work by 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 with those of a related competitive equilibrium and with

[16] This 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.

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

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 (1) each seller whose price strictly exceeds its cost sells its good, (2) each seller whose price is strictly less than its cost does not sell, and (3) the sellers from whom the auctioneer chooses to buy are the same as those who 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.[18] 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.[19]

The paper is organized as follows. Section II contrasts the performance of DA clock auctions and the Vickrey auction in a simple three-bidder example in which bidders may not be substitutes. Section III 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 IV characterizes the DA clock auctions as the set of auctions that preserve UWP. Section V introduces direct revelation DA auctions and shows that they are strategically equivalent to DA clock auctions. Section VI provides examples of DA auctions that may be useful in practice. Section VII 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, nonoptimizing one. Section VIII

[18] A stronger version of this equivalence obtains when the auction rule prescribes that 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 the 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. When 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.

[19] For 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.

compares the cost performance of DA auctions to two theoretical standards: competitive equilibrium and full-information equilibrium of paid-as-bid auctions. Section IX describes simulations of the performance of the FCC's reverse auction rules. Section X discusses multiminded bidders and their roles in the incentive auction. Section XI concludes.

II. A Simple Example

Here we illustrate properties of DA auctions with a simple example with three television stations (labeled 1–3) 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, and v_3. 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 acquire either 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 as 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").[20]

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

[20] The 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).

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; 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.

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.[21] 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 $[(w_1 + w_3)/w_2]v_2 \le v_2$, provided that $w_1 + w_3 \le w_2$, 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 toward minimizing the expected procurement cost in the style of Myerson (1981). Assume that bidder values are independently distributed with known distributions F_i, and

[21] In 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).

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 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.[22]

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.

III. 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

[22] In the actual FCC incentive auction, when the base clock price was q, the price offered to any feasible UHF television station i was $p_i(q) = q(w_i \cdot \text{Pop}_i)^{1/2}$, 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.

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.[23]

Informally, a 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 who 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.[24]

To avoid technical complications, we restrict attention to auctions with discrete time periods, indexed by $t = 1, \ldots$. 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: $A^t = (A_1, \ldots, A_t)$ such that $A_t \subseteq \ldots \subseteq A_1$. 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 $p : H \to \mathbb{R}^N$ such that $p(A^t) \leq p(A^{t-1})$ for all $t \geq 2$ and all A^t.

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 $A_{t+1} = A_t \backslash E_t$ 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 $I_i : H \to S_i$. 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

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

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

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.[25]

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. QED

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 nontruthfully 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. QED

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.[26]

IV. 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

[25] This 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.

[26] Like 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 et al. 2009). However, we would not expect bidders to play weakly dominated strategies if they have sufficient uncertainty about each other.

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. The protocol implements an allocation rule that can be described by $\alpha : \Pi_{i\in N} V_i \rightarrow 2^N$, 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.[27] We modify this definition in two ways. First, we weaken it by focusing only on winners' privacy. The DA auctions do not preserve losers' privacy, since their dropout 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 UWP if for any player i, any pair of his possible values $v_i, v_i' \in V_i$, and any values $v_{-\mathrm{i}} \in V_{-\mathrm{i}}$ of the other players such that the protocol results in bidder i winning in both states $(i \in \alpha(v_i, v_{-\mathrm{i}}) \cap \alpha(v_i', v_{-\mathrm{i}}))$, 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 $V = \Pi_{i\in N} V_i$.

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, player i must not exit in either case, and so the protocol must finish in the same terminal node in both cases.[28]

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

[27] This notion is also known as noncryptographic 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, Micali, and Lepinski 2005).

[28] This argument actually establishes a somewhat stronger property than UWP: that even if the other players use nontruthful 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.

direct mechanism for α is dominant strategy incentive compatible and therefore α must be monotonic (see sec. V). 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: $\bar{V} = \bar{V}_1 \times \ldots \times \bar{V}_N$. 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 $i \in \alpha(p_i, v_{-\mathrm{i}})$; by monotonicity of α we also have $i \in \alpha(v_i, v_{-\mathrm{i}})$ 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 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 (i.e., who have not exited the clock auction) must be winners. Thus, the constructed DA auction implements allocation rule α. QED

V. 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 $\alpha : \Pi_{i \in N} V_i \to 2^N$ and a payment rule $\pi : \Pi_{i \in N} V_i \to \mathbb{R}^N$ such that losing bidders are not paid; that is, $\pi_i(v) = 0$ for all $i \in N \backslash \alpha(v)$. 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:

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

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 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 $s_i^A : V_i \times V_{N \setminus A} \to \mathbb{R}_+$ is nondecreasing in its first argument. The input to the algorithm is a value vector $v \in V$, and the algorithm 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 that 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:

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

for every bidder $i \in A_{t+1}$. In the final round of the algorithm, for every winner $i \in A^T$, p_i^T 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 $N\backslash S$ never exits.[29] 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 the clock auction if he bids truthfully with value v_i and all bidders from $N\backslash A$ bid truthfully with values $v_{N\backslash A}$, while bidders in $A\backslash\{i\}$ never quit:

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

(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 ($\arg\max_{i\in A} s_i^A(v_i, v_{N\backslash A})$) 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 prices 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. QED

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

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

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 $i \in N\backslash A_t$, 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

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

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 becomes a winner, his final clock price is his highest value that would be still winning, that is, his threshold price. QED

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).[30] On the other hand, direct auctions lose obvious strategy-proofness and winners' privacy, making them less attractive for some applications.[31] 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.

VI. Examples of DA Auctions

In this section we describe some examples that shed light on potential applications. These examples expand on the examples of section II 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 feasible for the auctioneer to accept bids of only 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).[32] A DA clock auction maintains feasibility if it only reduces prices to stations i if $A \backslash \{i\} \in \mathcal{F}$ and only to one such station at a time.

In the reverse auction of the FCC's incentive auction, checking whether a given set $N \backslash A$ 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.[33] 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

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

[31] A 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, e.g., clock auctions with sealed intraround bids (restricted to be between start-of-round and end-of round prices) and the survival auctions of Fujishima, McAdams, and Shoham (1999).

[32] For instance, in the example in sec. II, $\mathcal{F} = \{\{1, 2\}, \{3\}, \{1, 3\}, \{2, 3\}, \{1, 2, 3\}\}$.

[33] In 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, Newman, and Leyton-Brown 2018).

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:

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

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: $V_i \cap V_j = \varnothing$ 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 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.[34]

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 II. 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

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

[34] This 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, i.e., 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 the 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.

The problem of maximizing $\Sigma_{i\in N} v_i$ subject to $\Sigma_{i\in N} w_i \leq 1$ 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 $\mathcal{F}$.[35] This formulation generalizes the example of section II, so we again appeal to the Dantzig greedy heuristic, 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 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 v_i/w_i multiplied by the remaining empty space in the knapsack when item i is first rejected.

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 II 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 $\mathbb{E}[\Sigma_{i\in\alpha(v)}\gamma_i(v_i)]$, 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 $\hat{s}_i^A(b_i, b_{N\setminus A}) = s_i^A(\gamma_i(b_i), (\gamma_j(b_j))_{j\in N\setminus A})$.

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 by Segal (2003), the auctioneer would optimally condition the price offered to one bidder on the bids of the others, including those who 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 $p_A^*(v_{N\setminus A}) = \arg\max_p(\pi - p)$

[35] In follow-up work, Duetting, Gkatzelis, and Roughgarden (2014a) 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.

$\Pr\{v_i \le p | v_{N \setminus A}\}$ 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 independent and identically distributed from an unknown distribution.

Another important objective in the 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 II, 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 $R(\varnothing) = 0$. 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 $\Sigma_{i \in A_t} p_i(A^{t-1}) > R(A_t)$.[36] 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, Roughgarden, and Sundararajan (2009). Their clock auctions always offer prices to all active bidders who exactly cover the cost of serving them (in our notation, $\Sigma_{i \in A_{t-1}} p_i(A^{t-1}) = R(A_{t-1})$ for all histories A^{t-1}) and stops as soon as all active bidders accept their prices (i.e., $A_t = A_{t-1}$). A different kind of budget-constrained (sealed-bid and clock) DA auction for procurement is proposed by Ensthaler and Giebe (2009, 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).[37] The FCC's incentive

[36] The 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.

[37] McAfee'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

auction is similarly a double auction for spectrum sellers (television 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.[38]

VII. Optimization, Substitutes, and Clock Auctions

A substantial literature on many-to-one matching (Gale and Shapley 1962; Kelso and Crawford 1982; Demange, Gale, and Sotomayor 1986; Hatfield and Milgrom 2005) has demonstrated that a stable allocation can be found using a DA 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 Ausubel 2004 and Bikhchandani et al. 2011 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 would not 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 substitutes property is satisfied in some classical examples (such as example 2), 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, or correlation in expected cost minimization). However,

---

demand within 1 of supply and stopping as soon as both (1) the sell price falls weakly below the buy price and (2) demand equals supply.

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

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

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

for some function $F : 2^N \to \mathbb{R}^N \cup \{-\infty\}$ and functions $\gamma_i : \mathbb{R} \to \mathbb{R}$. For example, condition (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, condition (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 $\alpha : \mathbb{R}_+^N \to 2^N$ solves condition (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 $V = \Pi_{i \in N} V_i \subseteq \mathbb{R}_+^N$ is implementable with a DA clock auction. Then, α has substitutes on any such state space V on which the objective in condition (3) has no ties.

The proof of the proposition is given in appendix D.

VIII. 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 (1) $v_i \le p_i$ for all $i \in A$, (2) $v_i \ge p_i$ for all $i \in N \backslash A$, and (3) $A = \alpha(p)$.

The following proposition establishes, for any DA allocation rule α, a three-way equivalence between (1) the winners' threshold prices $(\pi(v))_{i \in A}$ (which are also the DA auction prices), (2) the maximal competitive equilibrium prices supporting the winning allocation, and (3) 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. $(p_A, v_{N \backslash A})$ is a maximal price vector such that $(A, p_A, v_{N \backslash A})$ 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. Part b implies part a. We must have $A = \alpha(v) = \alpha(p_A, v_{N\setminus A})$ by equilibrium condition 3 and UWP. Thus, by equation (1), $p_i \leq \pi_i(v_{-i})$ for all $i \in A$. This must hold with equality for the equilibrium to be maximal.

Parts a and b imply part c. By part 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.

Part c implies part 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 2 and 3 by construction and satisfies condition 1 because of 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 2 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 $p'_A \leq p_A$; 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 because of UWP. QED

As noted above, 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 nonbossy, 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.

IX. 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 with those of the Vickrey-Clarke-Groves (VCG) mechanism. Since fully optimizing the repacking of the 2,990 US 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 US 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 very high frequency (VHF). This subset covers one of the densest areas of the United States: 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 (posted by the FCC at http://data.fcc.gov/download/incentive-auctions/Constraint_Files/). We simulated the problem of clearing 21 broadcast channels (126 MHz of spectrum), which was the FCC's initial clearing target in the incentive auction. Station 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 data set is $13.3 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 the FCC in the auction (posted at 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 more than 90 days of aggregate CPU time (the computation of VCG prices to the winners could be parallelized).

Then, we simulated the FCC's auction for the same data set, using the SATFC feasibility-checking software developed by Auctionomics (described by Frechette, Newman, and Leyton-Brown 2018). We used SATFC, version 2.3.1 (publicly released at https://github.com/FCC/SATFC/releases), which was also used by the FCC in the actual auction. In contrast, the average FCC auction on the scaled-down setting took 1.5 hours of aggregate CPU time.[39]

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%–30% lower than the costs paid by VCG (the average cost savings was 24%).

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

In addition, the FCC design proved to be computationally tractable and could be scaled up to the true size with a roughly proportional increase in computation time: indeed, a nationwide computation could be run using about 6 CPU hours.[40] 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.

X. Multiminded Bidders

Our analysis has restricted bidders to be single-minded, that is, to bid for a single predetermined option. The problem of algorithmic design of strategy-proof mechanisms for multiminded 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 multiround ascending auctions used by the FCC (Gul and Stacchetti 2000; Milgrom 2000), multiunit buyers have incentives to engage in demand reduction (and, similarly, multiunit 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, Cramton, and Milgrom 2006; de Vries, Schummer, and Vohra 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, Gonen, and Nisan 2003), randomized mechanisms that are truthful in expectation (Lavi and Swami 2011), and mechanisms that implement in undominated strategies (Babaioff, Lavi, and Pavlov 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 multiminded: 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,

[40] Even 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.

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.[41]

XI. 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 above. In contrast to sealed-bid auctions, DA clock auctions have the huge advantage for this application that 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.[42]

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

[41] Given the DA clock auction design, market power by broadcasters was a potential problem that 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.

[42] Two additional advantages make a DA clock auction particularly suitable for the FCC's problem and account for the 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.

as meeting certain other gross and net revenue goals). High costs could lead to less clearing, less trade, and a loss of efficiency.[43]

Roth (2002, 1341) has observed, "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 DA idea as the basis for some of the most successful new market mechanisms.

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." *Ann. Operations Res.* 153 (1): 79–129.

Ausubel, L. M. 2004. "An Efficient Ascending-Bid Auction for Multiple Objects." *A.E.R.* 94 (5): 1452–75.

Ausubel, L. M., P. Cramton, and P. Milgrom. 2006. "The Clock-Proxy Auction: A Practical Combinatorial Auction Design." In *Combinatorial Auctions*, edited by Peter Cramton, Yoav Shoham, and Richard Steinberg, 115–38. Cambridge, MA: MIT Press.

Ausubel, L. M., and P. Milgrom. 2002. "Ascending Auctions with Package Bidding." *Frontiers Theoretical Econ.* 1 (1): 1452–75.

———. 2006. "The Lovely but Lonely Vickrey Auction." In *Combinatorial Auctions*, edited by Peter Cramton, Yoav Shoham, and Richard Steinberg, chap. 1. Cambridge, MA: MIT Press.

Babaioff, M., and L. Blumrosen. 2008. "Computationally-Feasible Truthful Auctions for Convex Bundles." *Games and Econ. Behavior* 63 (2): 588–620.

Babaioff, M., R. Lavi, and E. Pavlov. 2009. "Single-Value Combinatorial Auctions and Algorithmic Implementation in Undominated Strategies." *J. ACM* 56 (1): 4.

Bartal, Y., R. Gonen, and N. Nisan. 2003. "Incentive Compatible Multi Unit Combinatorial Auctions." In *TARK '03: Proceedings of the 9th Conference on Theoretical Aspects of Rationality and Knowledge*, 72–87. New York: Assoc. Computing Machinery.

Bernheim, D., and M. Whinston. 1986. "Menu Auctions, Resource Allocation, and Economic Influence." *Q.J.E.* 101 (1): 1–31.

Bikhchandani, S., S. de Vries, J. Schummer, and R. V. Vohra. 2011. "An Ascending Vickrey Auction for Selling Bases of a Matroid." *Operations Res.* 59 (2): 400–13.

Blume, A., P. Heidhues, J. Lafky, J. Munster, and M. Zhang. 2009. "All Equilibria of the Multiunit Vickrey Auction." *Games and Econ. Behavior* 66:729–41.

Brandt, F., and T. Sandholm. 2005. "Unconditional Privacy in Social Choice." In *Proceedings of the 10th Conference on Theoretical Aspects of Rationality and Knowledge*, Singapore, June 10–12.

[43] We had also explored incorporating yardstick competition into the design, allowing the auction to set maximum prices for broadcasters in regions with little competition on the basis of bids in other more competitive regions. This option, however, was judged to be too complex.

Dantzig, G. 1957. “Discrete-Variable Extremum Problems.” *Operations Res.* 5:266–77.

Day, R. W., and P. Milgrom. 2008. “Core-Selecting Package Auctions.” *Internat. J. Game Theory* 36:393–407.

Demange, G., and D. Gale. 1985. “The Strategy Structure of Two-Sided Matching Markets.” *Econometrica* 53 (4): 873–88.

Demange, G., D. Gale, and M. Sotomayor. 1986. “Multi-Item Auctions.” *J.P.E.* 94 (4): 863–72.

de Vries, S., J. Schummer, and R. V. Vohra. 2007. “On Ascending Vickrey Auctions for Heterogeneous Objects.” *J. Econ. Theory* 132 (1): 95–118.

Doraszelski, U., K. Seim, M. Sinkinson, and P. Wang. 2016. “Ownership Concentration and Strategic Supply Reduction.” Discussion Paper no. DP11173, Center Econ. Policy Res., Washington, DC.

Duetting, P., V. Gkatzelis, and T. Roughgarden. 2014. “The Performance of Deferred-Acceptance Auctions.” In *Proceedings of the 15th ACM Conference on Economics and Computation*, Palo Alto, CA, June 8–12.

Duetting, P., T. Roughgarden, and I. Talgam-Cohen. 2014. “Modularity and Greed in Double Auctions.” In *Proceedings of the 15th ACM Conference on Economics and Computation*, Palo Alto, CA, June 8–12.

Ensthaler, L., and T. Giebe. 2009. “Subsidies, Knapsack Auctions and Dantzig’s Greedy Heuristic.” Discussion Paper no. 254. SFB/TR 15 (Sonderforschungsbereich/Transregio 15), Mannheim.

———. 2014. “A Dynamic Auction for Multi-Object Procurement under a Hard Budget Constraint.” *Res. Policy* 143:179–189.

FCC (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. 2018. “Deep Optimization for Spectrum Repacking.” *Communications ACM* 61 (1): 97–104.

Fujishima, Y., D. McAdams, and Y. Shoham. 1999. “Speeding Up Ascending-Bid Auctions.” In *Proceedings of the 16th International Journal Conference in Artificial Intelligence*, Stockholm, July 31–August 6, 554–59.

Gale, D., and L. S. Shapley. 1962. “College Admissions and the Stability of Marriage.” *American Math. Monthly* 69:9–14.

Gul, F., and E. Stacchetti. 2000. “The English Auction with Differentiated Commodities.” *J. Econ. Theory* 92 (1): 66–95.

Hatfield, J., and P. Milgrom. 2005. “Matching with Contracts.” *A.E.R.* 95 (4): 913–35.

Hayek, F. A. 1945. “The Use of Knowledge in Society.” *A.E.R.* 35:519–30.

Hurwicz, L. 1977. “On the Dimensional Requirements of Informationally Decentralized Pareto-Satisfactory Processes.” In *Studies in Resource Allocation Processes*, edited by K. J. Arrow and L. Hurwicz, 413–24. New York: Cambridge Univ. Press.

Izmalkov, S., S. Micali, and M. Lepinski. 2005. “Rational Secure Computation and Ideal Mechanism Design.” In *46th Annual IEEE Symposium on Foundations of Computer Science*, 585–94. Piscataway, NJ: IEEE.

Jarman, F., and V. Meisner. 2017. “Ex-Post Optimal Knapsack Procurement.” *J. Econ. 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*, edited by H. Liu, J. J. Salerno, and M. J. Young, 123–29. New York: 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* 55:1275–304.

Kelso, A., and V. Crawford. 1982. "Job Matching, Coalition Formation, and Gross Substitutes." *Econometrica* 50 (6): 1483–504.

Kim, A. 2015. "Welfare Maximization with Deferred Acceptance Auctions in Reallocation Problems." In *23rd European Symposium on Algorithms.*

Kushilevitz, E., and N. Nisan. 1997. *Communication Complexity.* Cambridge: Cambridge Univ. Press.

Lavi, R., and C. Swami. 2011. "Truthful and Near-Optimal Mechanism Design via Linear Programming." *J. ACM* 58 (6): 25.

Lehmann, D., L. I. O'Callaghan, and Y. Shoham. 2002. "Truth Revelation in Approximately Efficient Combinatorial Auctions." *J. ACM* 49 (5): 577–602.

Li, Shengwu. 2017. "Obviously Strategy-Proof Mechanisms." *A.E.R.* 107 (11): 3257–87.

Loertscher, S., and L. M. Marx. 2015. "Prior-Free, Privacy Preserving Bayesian Optimal Mechanisms." Working paper.

McAfee, R. P. 1992. "A Dominant Strategy Double Auction." *J. Econ. Theory* 56 (2): 434–50.

McMillan, J. 1994. "Selling Spectrum Rights." *J. Econ. Perspectives* 8 (3): 145–62.

Mehta, A., T. Roughgarden, and M. Sundararajan. 2009. "Beyond Moulin Mechanisms." *Games and Econ. Behavior* 67 (1): 125–55.

Milgrom, P. 2000. "Putting Auction Theory to Work: The Simultaneous Ascending Auction." *J.P.E.* 108 (2): 245–72.

———. 2009. "Assignment Messages and Exchanges." *American Econ. J. Microeconomics* 1 (2): 95–113.

———. 2017. *Discovering Prices: Auction Design in Markets with Complex Constraints.* New York: Columbia Univ. 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.

Milgrom, P., and I. Segal. 2017. "Designing the US Incentive Auction." In *Handbook of Spectrum Auction Design*, edited by Martin Bichler and Jacob Goeree, chap. 37. Cambridge: Cambridge Univ. Press.

Milgrom, P., and R. Weber. 1982. "A Theory of Auctions and Competitive Bidding." *Econometrica* 50 (5): 1089–122.

Moulin, H. 1999. "Incremental Cost Sharing: Characterization by Coalition Strategy-Proofness." *Social Choice and Welfare* 16:279–320.

Mount, K., and S. Reiter. 1974. "The Information Size of Message Spaces." *J. Econ. Theory* 28:1–18.

Mu'alem, A., and N. Nisan. 2008. "Truthful Approximation Mechanisms for Restricted Combinatorial Auctions." *Games and Econ. Behavior* 64 (2): 612–31.

Myerson, R. 1981. "Optimal Auction Design." *Math. Operations Res.* 6:58–73.

Newman, N., K. Leyton-Brown, P. Milgrom, and I. Segal. 2017. "Assessing Economic Outcomes in Simulated Reverse Clock Auctions for Radio Spectrum." Working paper, Cornell Univ.

Nisan, N., and A. Ronen. 1999. "Algorithmic Mechanism Design." In *Proceedings of the 31st Annual ACM Symposium on Theory of Computing*, 129–40. New York: Assoc. Computing Machinery.

Ostrovsky, M., and R. P. Leme. 2015. "Gross Substitutes and Endowed Assignment Valuations." *Theoretical Econ.* 10 (3) 853–65.

Oxley, James G. 1992. *Matroid Theory.* Oxford: Oxford Univ. Press.

Roth, A. E. 2002. "The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics." *Econometrica* 70 (4): 1341–78.

Segal, I. 2003. "An Optimal Pricing Mechanism with Unknown Demand." *A.E.R.* 93 (3): 509–29.

———. 2007. "The Communication Requirements of Social Choice Rules and Supporting Budget Sets." *J. Econ. Theory* 136:341–78.