Spectrum Auctions from the Perspective of Matching

Paul Milgrom and Andrew Vogt

May 5, 2021

1 Introduction

In July 1994, the United States Federal Communications Commission (FCC) conducted the first economist-designed auction for radio spectrum licenses. In the years since, governments worldwide have come to rely on auction mechanisms to allocate – and reallocate – rights to use electromagnetic frequencies. Over the same period, novel uses for spectrum have dramatically increased both the demand for licenses and auction prices, drawing continued attention to the nuances of spectrum markets and driving the development of spectrum auction design. In August 2017, the FCC completed the Broadcast Incentive Auction, a two-sided repurposing of an endogenously-determined quantity of spectrum that ranks among the most complex feats of economic engineering ever undertaken. The next generations of mobile telecommunications are poised to extend this growth, and to demand further innovation in the markets and algorithms that are used to assign spectrum licenses.

What does all of this have to do with matching theory? Spectrum auctions differ from canonical matching problems in meaningful ways. Market designers often emphasize that a key element of matching markets is the presence of preferences on both sides of the market. In a marriage, for example, it is not enough that you choose your spouse; your spouse must also choose you. This two-sided choice structure applies also to matches between students and schools and between firms and workers, but not to matches between telecommunications companies and radio spectrum licenses.

It is a different matching element that is often critically important in radio spectrum auctions. A radio spectrum license grants its holder the right to use a particular band of frequencies in a particular geographical area, subject to restrictions on signal power, interference with nearby uses, and other regulations. Radio spectrum, like most physical resources, is heterogeneous. Signals sent using different frequencies travel differently. Generally, lower frequencies can support communication over greater distances, passing through diverse obstructions like raindrops, trees, and concrete walls. That reach is an advantage for serving large rural areas or for providing outdoor coverage among tall buildings. Higher frequencies are less useful for long-distance transmissions, but their short range is an advantage for supporting multiple users simultaneously, as when neighboring households share the same frequencies for their Wi-Fi applications. Even very small differences in frequencies can be important, as interference can be caused by harmonics between frequencies being used in the same device. An efficient allocation must match the right sets of spectrum licenses to the right users: that is the critical matching element that connects spectrum auctions to matching markets.

The economic analysis of radio spectrum auctions also addresses complications of license assignment that are less common in other matching markets. The theory of spectrum auctions borrows from textbook market theory, because efficient spectrum allocations must also assign the right *quantities* (amounts of bandwidth) to different agents. The need to ration quantities efficiently leads naturally to a central role for prices. Although the matching literature includes models of matching with prices, it is difficult to implement those models using direct mechanisms, which require participants to report their complete preferences, because multi-product demand functions based on multiple prices embed much more information than simple rank order lists. The theoretical connection between matching and auctions typically assumes away this difficulty to expose similarities between matching with and without prices.

Rather than requiring buyers to exhaustively describe their demands, spectrum auctions instead use dynamic mechanisms, in which bidders can respond in more complicated ways to incrementally evolving prices. These dynamic auctions provide some advantages, allowing bidders to express a wider set of preferences and accommodate budget constraints, but they also alter bidding incentives and, in particular, weaken the incentives for truthful reporting of values. The relative magnitudes of these advantages and disadvantages vary across applications, which helps to account for the historical successes and failures of certain spectrum auctions.

Most of the simple, elegant results of matching theory assume that agents report preferences that satisfy a substitutes condition, so another important concern in auction design is the failure of that condition, which can also be described as the presence of complementarities (see chapter **??**). Bidders in spectrum auctions often view licenses as complements: buying one frequency increases the value of another. This can arise when a bidder prefers to hold licenses for adjacent frequencies or licenses that cover adjacent territories, when there are returns to operational scale and scope, or when budget constraints force a choice between one expensive license and several cheaper ones. The seriousness of the problems that result from the failure of the substitutes condition and the best practical remedies to address them vary depending on the source of the complementarities.

Another important difference between traditional matching theory and auction theory lies in the welfare analysis. In most countries, the spectrum regulator is less concerned about the profits of the telecommunications companies, or even about expected revenue, than about consumer welfare and how the auction outcome will affect competition in the consumer market. The regulator often prefers a market in which many companies compete to provide consumers with low-cost, high-quality services that are widely available, even in small communities in rural or mountainous areas where the private cost of providing service may be high. It is not possible to understand the market design decisions made by spectrum regulators around the world without understanding the salience of their concerns about consumer welfare, retail competition, and universal service.

This chapter begins, in Section 2, by reviewing the deep connections between two-sided matching markets with and without money. The extension of the two-sided matching problem to incorporate prices subsumes, as special cases, not only models of matching without prices and of bilateral matching with prices, but also certain auction problems. We then describe the processes of early spectrum auction designs and their properties under naive bidding. The first spectrum auction designs were adopted after theorists demonstrated that, under certain assumptions about bidders' values (substitutes) and behaviors (naive), those designs could lead to allocations and prices approximating a competitive equilibrium.

In Section 3, we turn to the issue of incentives in spectrum auctions. Most often, buyers in these auctions seek to acquire multiple licenses, so matching theorems about dominant strategies for buyers of a single item do not apply. The most important competition issues for market design are similar to those present in homogeneous goods markets, such as anti-competitive behavior, market concentration, and broader societal interests.

Section 4 focuses on complementarities. Although the licenses sold in spectrum auctions are usually at least partial substitutes for one another, it is also common for bidders to view some licenses as complements. This is at odds with the early matching and auction literature, which mostly assumed that the items to be allocated are substitutes, or that the markets are large in ways that make failures of substitution unimportant. Radio spectrum auctions rarely have many competing participants, and complements in radio spectrum create meaningful complications, which have led to the adoption of new designs and to interesting subtleties in auction bidding strategies.

Section 5 explores the theory of descending clock auctions for a wider range of problems than just those in which items are substitutes, with particular attention to the problem of designing the Broadcast Incentive Auction. Descending clock auctions are shown to be a new kind of Deferred Acceptance Algorithm, and to share properties with DA Algorithms in the substitutes case. We describe seven significant disadvantages of the classic Vickrey auction design, each of which can be avoided by this class of clock auctions. Section 6 concludes, and describes other contexts to which similar mechanisms may be applied in the future.

2 Spectrum Auction Algorithms

2.1 Static Matching with Prices

Since the introduction of models for matching without prices (the Gale-Shapley college admissions or marriage problem) and one-to-one matching with prices (the Shapley-Shubik house assignment game), it has been clear that the two theories are closely connected. Those foundational matching models can be combined and generalized within the Kelso-Crawford model of matching between firms and workers: each worker can work for only one firm, but each firm can employ multiple workers, with wages selected from a finite set. Both firms and workers have preferences over their match partners and the contractual wage. With the important assumption that workers are substitutes, such that increasing the wage of one worker does not reduce a firm's demand for any other worker, a deferred acceptance process with offers that include prices leads to a stable allocation.

When the set of wages is a singleton, the firm-worker model is logically equivalent to the college admissions model, but with a more general substitutes preference structure. If each worker always prefers the firm offering her the highest wage, the Deferred Acceptance Algorithm in this model amounts to a simultaneous auction in which workers offer their services to the highest bidders. If, instead of using a direct mechanism, this process is run dynamically with automated bidders that represent the firms, then the model anticipates many of the rules of the simultaneous multiple round auction later used for selling radio spectrum licenses. If each firm has only a single opening, the matching problem can also be reduced to the model of one-to-one matching with prices. Thus, the firm-worker model nests the earlier models of matching with and without prices, and allows the reinterpretation of the Deferred Acceptance Algorithm as an auction mechanism with automated bidders.

The substitutes condition provides an important simplification in this matching model, because it is exactly the condition that ensures that during the labor-market auction with rising wages, a firm never chooses to reject a worker that it had wanted to hire when the wages of other workers were lower. Consequently, the demand for a given worker can only fall when that worker's own wage rises, which happens only when there is excess demand. Wages in the auction continue to rise until demand for each worker is exactly one, which is market-clearing.[1]

Connections between matching problems and spectrum auctions can also be found within the (further) generalized model of two-sided matching with contracts (see chapter **??**). The intuition is again that an auction is equivalent to a particular matching of buyers to sellers, with the restrictions that buyers

[1]Although this intuition is nearly correct, there are some challenging technical details to be addressed because, at some wage vectors, firms may be indifferent among workers, so the firms' demands functions may not be well defined and the formal analysis must use demand correspondences. These details, however, can be managed to find an exact equilibrium. With the substitutes condition, ascending auctions with discrete increments will also find approximate market equilibria. When the substitutes conditions fails, however, new issues can arise, as described in Section 4.

are differentiated only by the price they are willing to pay and that sellers care only about those prices. If there is only one firm or seller – an auctioneer – then the model of matching with contracts can be reduced to a particular auction mechanism called the generalized clock-proxy algorithm. In a clock-proxy auction, as in other deferred acceptance matching algorithms, the offering side must express its preferences using an exhaustive and static rank order list. However, many of the auctions designs that have been used to allocate spectrum rights in practice instead employ a dynamic process in which the offering side can incrementally adjust its bids.

2.2 Simultaneous Multiple Round Auction

The inaugural 1994 FCC spectrum auction used the newly-invented simultaneous multiple round (SMR) auction format. This design has since become common for government sales of spectrum licenses around the world. It uses a dynamic process that resembles the firm-proposing Deferred Acceptance Algorithm with prices, in which the radio spectrum licenses to be sold correspond to workers who each prefer the highest money offer they receive.

In the SMR auction, multiple items are offered for sale by a single auctioneer. Participants bid in a series of rounds, in which each new bid for an item must be submitted at a higher price, subject to a minimum increment. Each round of bidding is conceptually similar to a round of tentative offers in the DA Algorithm. The notion of a best proposal is replaced with that of the "standing high bid" (and its associated "standing high bidder"), which is the highest bid submitted for the item to that point in the auction (and the bidder that submitted it). Similar to the process in the DA Algorithm, bidding continues until there is a round in which there are no new bids for any item. At that point, the standing high bids become winning, which determines both the allocation and the final prices.

Mechanism 2.1 (*Simultaneous Multiple Round Auction*).

1. **Initialization**: The auctioneer sets a reserve price for each item. The initial standing high bidder for each item is the auctioneer, and the initial standing high bid is the reserve price.

Until there is a round in which no new bids are made, do:

2. **New bids**: Bidders submit new bids for any desired items, indicating the prices they are willing to pay. Each new bid for an item must be no less than its minimum allowable price.

3. For each item, the auctioneer tentatively accepts a standing high bid, equal to the maximum of the previous standing high bid and any new bids. Ties are broken randomly. The auctioneer reveals the standing high bid and standing high bidder for each item.

4. The minimum allowable price for each item is set equal to its standing high bid plus an additional increment.

Output the matching of items to bidders and the prices for each item.

There are clear similarities between the DA Algorithm and the SMR auction, but there are also important differences. The DA Algorithm is a direct mechanism, in which participants report what they would demand from any available set of items. With multiple items and prices, specifying a complete demand function can be a daunting task. The SMR auction, by contrast, is a dynamic mechanism in which bidders are never required to make exhaustive reports. Instead, they receive information about each round's results, and can use that information to adjust their next bids. This eliminates some of the guesswork required under static formats (such as sealed-bid auctions) and lowers the risk of ex-post bidder regret. Another effect of this dynamic reporting regime is to allow bidders to express some demands that are incompatible with substitutes preferences; perhaps most importantly, bidders can ensure that their bids across all items respect overall budget constraints.

However, the rich set of bidding possibilities within a dynamic mechanism also enables strategies that may be harmful to the operation of the auction. One example arises when a bidder limits its participation early in the auction, because it prefers to wait to see how high the prices for one item may rise before it decides how much to bid for another. If many bidders refrain from bidding early in the auction, with each attempting to be the last to bid, then some bidders are sure to be disappointed, and the number of auction rounds could grow impractically large. The FCC's original SMR auction design mitigated that problem by adopting the "Milgrom-Wilson *activity rule*," which limited each bidder's bids in any round relative to its overall bidding activity in the previous round. All subsequent SMR auctions and related dynamic designs have incorporated some sort of activity rule.

The outcome of an SMR auction resembles that of the DA Algorithm if participants bid as if items were substitutes. As in ordinary demand theory, items are said to be *substitutes* if, whenever a particular set of items uniquely maximizes a bidder's profit at some prices, that same set remains a profit-maximizing choice at any new price vector with the same prices for items that were being demanded and weakly higher prices for all other items. When each bidder in an SMR auction has preferences satisfying the substitutes condition and naively regards the prices it faces as fixed – equal to the current price for items for which it is the standing high bidder and to the current price plus one increment for all other items – the SMR auction exactly follows the rounds of the DA Algorithm. In this case, the final prices and outcome form a competitive equilibrium of an economy that is "close" to the actual one, in which participants' values for each combination of items are reduced by one increment for each of the items in that combination that the bidder does not win in the auction. This implies that any inefficiency is at most proportional to the size of the auction bid increment.

2.3 Clock Auction with Assignment Round

In many recent spectrum auctions, licenses are grouped into product categories within which the licenses are mostly interchangeable. In that case, the SMR auction can be closely approximated by a clock auction in which a single price applies to all licenses in the same category. For example, the regulator might auction licenses with slightly different frequencies that cover the same geographic area, and group these together in a single product category. Given the prices, each bidder specifies the *quantities* of each product that it wishes to buy. For each category in which licenses are overdemanded in the current round, the auctioneer raises the prices for the next round. Other prices remain unchanged. A bidder may increase but may not reduce its demand for a product with an unchanged price.[2] This auction process continues until there is no excess demand for any product category.[3]

A clock auction mechanism determines the quantities of each license category won by each bidder, but not the assignment of specific licenses. In spectrum auctions, bidders often care about the particular frequencies they will be allowed to use. Efficiency typically demands that all winning bidders should receive contiguous blocks of spectrum, to minimize interference issues at frequency boundaries with different licensees. Perhaps all bidders also wish to avoid a particular "impaired" frequency that conflicts with existing uses of the same or nearby

[2]Notice that if the different products are substitutes for a bidder, then it has no wish to reduce demand for licenses whose prices have not increased.

[3]Standard clock auctions do not use the notion of a standing high bidder, but this creates the possibility that in some round, demand may start strictly above the available supply but fall strictly below during the round. Most clock auctions implement additional rules to avoid this, refusing to honor demand reductions that would result in demand for a product that is less than supply.

frequencies. Perhaps one bidder wants to avoid a particular frequency because of harmonic interactions with another of its licensed frequencies.

To determine the final license assignment and address these issues, the clock auction is followed by a final bidding round called the "assignment round." In a typical assignment round, each bidder submits a set of bids that express its values for alternative assignments, given the quantities of licenses that it has won in each product category. The auctioneer selects the bids that maximize the total value, usually subject to the constraints that each winner is assigned a contiguous block of frequencies of the correct size and that no two assigned blocks intersect. These become the winning bids, and the additional payments to be made in the assignment round may be set by a first or second price payment rule.

Mechanism 2.2 (*Clock Auction with Assignment Round*)**.**

1. **Initialization**: The auctioneer sets an initial clock price for each product category.

Until there is a round in which no new bids are made, do:

2. **New bids**: Bidders submit new bids for any desired products, indicating their quantity demand at the current clock prices.
3. The auctioneer reveals the aggregate demand for each product. The clock price for each product that is over-demanded is set equal to its previous price plus an additional increment.

As an intermediate step, output the matching of products to bidders and the prices for each product.

4. **Assignment Round**: Accept additional bids to assign bidders to specific items, consistent with the matching of products to bidders.

Output the matching of items to bidders and the prices for each item.

Because clock auctions aggregate similar items into larger product categories, they simplify bidding and speed up auctions. For example, suppose there is a category of ten nearly identical spectrum licenses, and eleven bidders that each want only one. Bidders can (nearly) describe their preferences with just one bid price, rather than ten, which simplifies bidding. The auction process is also faster, because an SMR auction would require ten rounds with different standing high bidders to raise the price for all items, with the price of only one license rising by one increment in each round. In contrast, a clock auction that categorized the licenses as identical would raise the single price associated with all ten licenses in one round. In this way, this clock auction requires both fewer bids in each round and fewer rounds to reach completion.

3 Bidder Incentives and Regulator Objectives

3.1 Strategy-proofness in Matching and Auctions

Matching theory and auction theory also each include results about dominant strategy incentives that are closely connected. In the one-to-one matching problem, it is a dominant strategy for the members of the offering side to report their preference lists truthfully. This finding is closely connected to the application of the Vickrey auction within the Shapley-Shubik assignment model for "houses" (indivisible goods with unit supply and demand), which yields similar strategy-proofness.

To understand this connection, it is best to keep firmly in mind that participants' incentives in any direct mechanism depend only on the mapping from reports to outcomes, and not on the details of the algorithm used to compute the outcome. For the one-to-one matching problem, it is not the deferred acceptance process itself that is important, but that the mechanism computes the most favorable stable assignment for one side. In the house allocation problem, if we identify the value-maximizing allocation and set Vickrey prices for the buyers, then each buyer pays its opportunity cost for the house it acquires. That vector of house prices is the lowest market-clearing price vector, and thus is the buyer-preferred stable outcome for the housing market. For both the one-to-one matching problem and the house allocation problem, the mechanisms that select the most favorable stable outcome for one side are strategy-proof for that side. In one case, the mechanism is usually described in terms of the DA Algorithm; in the other, it is described as a Vickrey auction. Both are special cases of the same mechanism: they map reported preferences to the most preferred stable outcome for one side of the market, and therefore they provide that side with an incentive to report truthfully.

The DA Algorithm also selects the most preferred stable outcome for one side in the general model of many-to-one matching with contracts (which may include prices and other terms), provided that the side that accepts multiple contracts has substitutes preferences that satisfy the Law of Aggregate Demand, which requires that those participants select weakly more contracts when selecting from an expanded set of alternatives. If the mechanism selects the most preferred stable outcome for the single-contract side, then participants on that side maximize their values by reporting truthfully. These truthful-reporting incentives can be established using only the fact that the mechanism chooses the most-preferred stable outcome. For the related model in which a firm wants to hire multiple workers, the vector of the workers' Vickrey wages is the highest competitive equilibrium price vector, so it is the worker-best stable allocation. Consequently, workers have incentives to report truthfully.

3.2 Market Power Concerns

The dominant strategy theorems of matching theory and the related results in auction theory all describe incentives for truthful reporting that apply just to agents that want to be matched to only one other agent or to buy only a single item. For most spectrum auctions, this condition does not hold, which means that in practice bidders have predictable incentives to misreport their values. When the number of bidders is also small, concerns about their market power take on increased importance.

In ordinary textbook markets, producers exercise market power by restricting output to generate higher prices. In spectrum auctions, large bidders can sometimes exercise market power using a strategy of "demand reduction": by demanding fewer licenses, the bidder causes the auction to clear at lower prices. Spectrum auctions sometimes attract only a small number of serious bidders, each an incumbent telecommunications company with the resources to estimate the plausible preferences and license values of its opponents. In that setting, the results of demand reduction strategies can be dramatic. Several spectrum auctions have failed after the major bidders all reduced demand in an early round, producing substantially lower spectrum prices relative to other similar auction. A German auction in 1999 concluded after only two rounds of bidding, with two large bidders splitting the available spectrum 50-50. Ofcom's 2021 auction of spectrum in the 700 MHz and 3.6 GHz frequencies ended at low prices with an equal division of spectrum among the winning bidders in each band.

This behavior is hardly surprising. Consider a clock auction under a simplified setting, in which there is just one category of items available and bidders have marginal valuations that are decreasing in the number of items acquired. Assuming complete information and small price increments, the unique equilibrium strategy that survives iterated elimination of weakly dominated strategies is for all winning bidders to forego competition and bid from the outset for their final quantities. Prices greater than the reserve can arise in a such an equilibrium only if the auction includes losing bidders who win nothing.

The role of dynamic multiple-round auctions in facilitating these low-price outcomes is a subtle one. Sealed-bid, uniform-price Treasury bond auctions are, in principle, vulnerable to a similar kind of manipulation. With complete information, those auctions have an equilibrium with high bids for just the quantities bidders expect to buy in the auction and low prices or no demand for any higher quantity. The auction market then clears at the reserve price. Only bidders who win zero in equilibrium bid for more than they will win, and those bidders set the clearing price.

3.3 Small Numbers and Regulator Objectives

Even setting aside the matching element, there are two things that distinguish spectrum auctions from their Treasury auction counterparts. One is the very small number of bidders and the absence of uncertainty created by other bidders. With two bidders and an even number of spectrum licenses, for example, there is a focal outcome in which each bidder bids for half of the available spectrum. The second relates to how bidders settle on an expected final outcome. Although there is currently no theory of communication in dynamic auctions, bidders appear sometimes to be able to coordinate their expectations through a series of bids, settling on a market division in which the winners avoid competing against one another in ways that might drive up auction prices. In other cases, where the likely preferences of competing bidders are common knowledge, strong bidders may be able to unilaterally enforce particular settlements by threatening to raise prices on licenses that their opponents desire.

In contrast with labor, school or medical matching, the regulator designing a spectrum auction is most often a government regulator that cares not only about revenues, but also about reducing concentration in the retail market in which the spectrum buyers compete. Buyers calculate values in spectrum auctions by forecasting revenue and cost streams with and without the spectrum and then computing the difference, so part of the value of any license comes from denying that spectrum resource to a competitor, which typically leads to an increase in the bidder's market share. Since the value of this market foreclosure (limiting the opportunities available to competitors) are usually higher for dominant firms, allocations that maximize these values lead to too much foreclosure of weaker firms. In some cases new entrants have been reluctant to compete for licenses at all, given the high costs of auction preparation and the likelihood that established bidders will bid above the use value of the spectrum, as measured by cost avoidance, in order to maintain their larger market positions.

Regulators also generally care about promoting better service to rural and mountainous areas where service is costly to provide. Good rural phone service makes it easier to provide emergency services and to keep citizens informed and engaged, among other public benefits.

In practice, regulators advance these interests and others in multiple ways, some of which involve the auction rules. To block foreclosure, regulators may use spectrum caps, which limit the quantity of spectrum that can be won by any single bidder; set-asides, which disqualify some incumbent firms bidders from bidding on certain licenses; or discounts for new entrants, which require small bidders to pay only a fraction of the value of their winning bids in an auction. To promote rural phone service and to further limit foreclosure, licenses may come with build-out requirements, such that the winner is required to return the license if it fails to build supporting infrastructure to use the spectrum over a sufficiently large area, or over rural areas in particular. The efficiency of any individual auction is almost always desirable, but from a regulatory perspective may be only one of several competing interests in the overall market.

4 Substitutes and Complements

The earliest models of matching theory and early auction models share the common assumption that each participant – firm or worker, buyer or seller – is looking for only one partner or only one good. Goods cannot be complements in those one-to-one matching models.

When the model is expanded to allow one side to desire multiple matches – say a firm that wants to hire several workers – the simplest generalizations involve assuming that goods or people are substitutes. The Gale-Shapley DA Algorithm finds a stable outcome for the college admissions problem when colleges regard students as substitutes but not generally otherwise, and a similar finding applies to the Kelso-Crawford model when firms regard workers as substitutes. It is not just that the algorithms themselves fail: without substitutes, stable matches may not exist. One example arises when two doctors in a couple want jobs in the same city and so regard offers from hospitals in any one city as complements: the offers are valuable as a pair, but of little value individually. Similar results apply in equilibrium theory. When goods are substitutes, the SMR auction under straightforward bidding finds a near-equilibrium and the ascending clock auction finds an exact competitive equilibrium, but both can fail when goods are not substitutes. In an auction market with a single seller and many buyers, the Vickrey outcome is in the core of the market game if goods are substitutes and there are ascending auctions that mimic the Vickrey outcome. Without substitutes, Vickrey outcome is not generally in the core and prices can be too low to be market clearing.

Just as the assumption of substitutes is inconsistent with a pair of doctors seeking residencies at hospitals in the same city, it is also at odds with observed bidder demand for spectrum licenses. Spectrum licenses can be complementary for a variety of reasons. Users of radio spectrum often have greater value for adjacent frequencies than for slightly separated ones, because adjacency reduces the number of boundary frequencies at which users need to coordinate with others to avoid interference. Users also prefer to license the same frequencies in adjacent geographic areas to avoid the need to coordinate their uses with others at the geographic borders. Similarly, there can be significant economies of scale and scope to make packages of licenses more valuable than the sum of individual values: licenses to use tiny amounts of frequency bandwidth or small geographic coverage areas may even have zero value because they must be combined with others to merit the infrastructure investments needed for a viable wireless system. Bidder budgets, too, can make licenses into complements. For example, suppose that a buyer has values of $(30, 30, 40)$ for three items A, B and C, and sets of licenses are valued at the sum of the item values. With a budget exceeding \$100, the demands for each license would depend just on its own price. If the budget is just \$10, however, and the price vector is $(5, 5, 10)$, then the buyer maximizes its net value by demanding A and B, but if the price of A then rises by 1, then the buyer maximizes by dropping both A and B to demand C alone. That pattern, in which a price increase for A reduces the demand for B, is what it means for the two licenses to be complements in demand, at least over that range of prices.

It is no coincidence that the substitutes condition so often fails in practice, because the condition is non-generic: substitutes valuations must satisfy a potentially large collection of linear equations. Given a set of items $S = \{A, B, C, \dots\}$, any three elements A, B, C, and any subset $T \subseteq S - \{A, B, C\}$, if v is a substitutes valuation, then the two largest numbers from the set

$\{v(AT) + v(BCT), v(BT) + v(ACT), v(CT) + v(ABT)\}$ must be equal.[4] With many goods, this leads to a large set of equations that v must satisfy. If S is a large set, this implies many restrictions on v, so the class of substitutes valuations is of much lower dimension than the unrestricted class of valuations.

An important way in which spectrum auction design falls outside the scope of matching market design is the way the spectrum products are defined. In standard matching theory and the markets to which it has most often been applied, the items to be matched (people, jobs, organs, online ad impressions) are defined *a priori*, but the same is not true for spectrum licenses. In a hypothetical world, one might imagine that regulators would prefer to sell 'postage stamp' licenses with tiny geographic areas and paper-thin bands of spectrum, allowing bidders to assemble bespoke collections of these small licenses to meet their operating needs. But such licenses are hardly likely to be substitutes, and trying to sell them using either a matching-like mechanism or an SMR design has no good theoretical properties and risks dramatically inefficient outcomes. To mitigate this problem most effectively, the regulator must carefully define the bandwidth, geographic coverage, time period, technological restrictions and regulatory restrictions that apply to the licenses.

In many countries, mobile operators operate nationwide networks, so there is little to be gained by offering spectrum licenses that serve smaller geographic areas. For this reason, many regulators commonly sell only nationwide licenses, which eliminates the geographic adjacency problem, fully addressing that kind of complementarity. For these same countries, if a new frontier of contiguous frequencies is offered for sale, then an assignment round following a clock auction can also guarantee the licensing of contiguous frequencies to all the winners. In more recent auctions, where old uses are being shut down to make room for new ones, frequencies may not be contiguous and other bidders may have already licensed some of the intervening frequency blocks. In those settings, to minimize the frequency contiguity problem, design often involves offering incentives to existing licensees to include their old blocks in the auction.

4.1 Exposure Risk

Bidders often differ in their demands in ways that make it impossible to package license rights to reflect the interests of every buyer. In such cases, one approach is to tailor license definitions to serve the needs of a particular class of bidders. For example, FCC Auction 66 simultaneously allocated licenses that covered the entire United States using three different geographic scales: six large Regional Economic Areas (REAs), 176 Economic Areas (EAs), and 734 Cellular Market Areas (CMAs). The CMAs served the needs of existing small cellular service providers, allowing them to buy additional bandwidth for the areas already served. The EAs and REAs suited the requirements of rural operators

[4] Suppose that, for given valuations, the first of these terms is larger than the other two. Then there exists some price evolution such that increasing the price of B can reduce demand for C, as the bidder switches from bidding for BC to bidding just for A.

and others serving areas larger than the CMAs but smaller than the national operators.

The alternative is to allow bidders to assemble collections of smaller licenses that, taken together, meet their business objectives. This can work well for an SMR auction when licenses are substitutes, because if a bidder bids straightforwardly according to its true demand, then a standing high bidder for a license after any auction round will still want to buy that license even if subsequent competition raises the prices of other licenses. Without the substitutes restriction, however, a straightforward bidder can come to regret its past bids. A buyer who wants and bids for the package AB may become the standing high bidder on the license A, but subsequent bidding may increase the price of license B. The substitutes condition asserts exactly that such a price increase does not reduce the bidder's demand for A, so it will never wish to withdraw its demand for A.

If valuations do not satisfy the substitutes condition, then if the bidder bids for A during the auction, it is exposed to the risk of winning that license when it wishes it had not. If it does not bid for A, it risks missing a profitable opportunity. This challenge to bidding well in an auction is called the *exposure problem*, and it has become a central concern in modern spectrum auction design.

A bidder in an SMR may be tempted to respond to the exposure problem by sitting out the early part of the auction, waiting to bid until late in the auction when it may have a better idea of what the prices for different items will be. For example, the bidder who wants package AB and is uncertain about whether B will be too expensive might choose to bid for B only and, upon learning that it can win that at a low price, might only then begin to bid for A. If many bidders behave that way, then the auction could stall – after all, it is not possible for every bidder to be the last to bid in an SMR auction. Such behaviors are limited by activity rules, but those rules cannot shield bidders from the risk of acquiring a less valuable subset of their desired licenses.

4.2 How Bidders Manage Exposure Risk

The theoretical treatments of the exposure problem focus on the possibility that a bidder without substitutes preferences can be caught on the horns of a dilemma: the bidder must either curtail its participation or risk losing money. Closer analysis reveals, however, that the magnitude of the exposure problem depends on price uncertainty. In the example above, if the bidder could narrow its forecast for the price of license B, it would find that one of its two main options – bidding aggressively or dropping out early – would almost always lead to a near-optimal outcome. Generally, if a bidder expects that, even without its own participation, auction prices will likely be very high or very low, then its exposure risk is easily managed.

Even when a simple high-versus-low characterization does not apply, there can be strategies that substantially mitigate the exposure problem. One notable example arose in FCC Auction 66, which used an SMR auction design to allocate "Advanced Wireless Service" (AWS) licenses. Auction 66 included a bidder called *SpectrumCo* – a consortium of cable operators with a business plan calling for it to acquire licenses covering all the major cities of the United States where cable services were offered. SpectrumCo's bidding team was initially instructed that if SpectrumCo could not acquire national coverage, it should not buy any coverage at all – codifying an extreme exposure risk and creating an unrealistic objective for the team. The operators eventually agreed that spectrum could be acquired so long as it was sufficiently cheap by historical standards, as this would allow the consortium to resell any licences it acquired without incurring large losses.

The main remaining risk for SpectrumCo came from the details of the auction design. The FCC set the opening price for each license in proportion to its bandwidth and population coverage, and specified the same percentage bid increments for every license. In practice, however, license prices are almost always much higher even in *per capita* terms in densely populated areas, where spectrum bandwidth per user is in shorter supply. Prices in the REAs of the Western States and Northeast were expected to be at least twice as high, relative to population and bandwidth, as prices in the Mississippi Valley and the Mountain States. If bidding proceeded with bids only at the minimum increments and SpectrumCo bid straightforwardly, it could become the standing high bidder committed to buy licenses in those low-price areas before it learned much about demand and prices for the high-value REAs.

SpectrumCo dealt with that problem by making the largest "jump bid" in the history of spectrum auctions – an increase in its total bid across all REAs of almost $750 million! Its bid doubled the prices of the valuable REAs to bring them closer to the clearing prices and to improve its estimation of the likely demand for those licenses.[5] When the demand of others bidders collapsed on the day of the jump bid, the SpectrumCo bidding team was able to forecast that its budget would enable it to acquire nationwide coverage, and SpectrumCo ultimately won licenses across the country.

The lesson here is a subtle one. On one hand, bidders in spectrum auctions are not helpless automata; they are strategic participants who can sometimes adopt strategies that mitigate the exposure problem. Doing so involves using information that emerges during the auction (especially about demands at different price levels) that is not revealed by any traditional direct mechanism. Still, these examples are all special, and the exposure problem is a real concern. This has inspired attempts to create mechanisms that avoid the exposure problem more generally, without the need for special adaptations by bidders.

4.3 Combinatorial Auctions

It is possible, in principle, to design a spectrum auction to avoid the exposure problem completely. Bidders in such *combinatorial auctions* submit all-or-nothing *package bids* for collections of licenses. The best-known combinatorial auction design is the *Vickrey auction*ch:spectrumauctions:vickreyauction,

[5]The jump bid also had a strategic justification, which we will not discuss here.

in which a bidder expresses its values for every desired combination of licenses. Given the bids, the auctioneer computes the assignment that maximizes the total value of the license assignment and then sets prices according to the famous Vickrey formula, in which each bidder pays the opportunity cost of the licenses it acquires – that is, the difference between the maximum value of all the licenses to the other bidders and the value ultimately assigned to the other bidders.

Vickrey auctions are challenged in several ways when licenses are not substitutes; problems include low prices and vulnerability to collusion even by losing bidders. Corresponding first-price auctions avoid these problems in selected full information equilibria, but there are also inefficient equilibria in that setting, and there are no general analyses of the incomplete information case.

A serious concern about both Vickrey and first-price auction designs is that their good theoretical properties depend on bidders being able to express values for every combination of licenses that they might win. With n distinct licenses, there are $2^n - 1$ such combinations, which for many spectrum auctions is an impractically large number for bidders to evaluate. These sealed-bid designs are most valuable when the number of packages of licenses that a bidder can win is small.[6]

Various designs for dynamic package bidding have been suggested, but the multi-round package auctions that have most often been used for spectrum sales are variants of the combinatorial clock auction (CCA). The CCA format has multiple stages. Bidders initially participate in a clock auction, in which the auctioneer announces prices for categories of items. In each round, the price is increased by one increment for each category with excess demand, until there is no excess demand for any product. Unlike a standard clock auction, a bidder's demand in each round of a CCA is understood as a package bid, which means that the bidder offers to buy all the items in its package or none, at a total cost not exceeding the current price of its package. It does not offer any price to buy a subset. The clock auction stage of the CCA ends when there is no excess demand for any product.

If bidders were to bid their true demands straightforwardly and if bidder valuations satisfy the substitutes condition, then this procedure leads to the same outcome as a clock auction. That fact provides some connection between the CCA format, clock auctions, and the Kelso-Crawford matching theory.

When licenses are not substitutes, no formal theory justifies this procedure, but the prices may help guide bidders to identify relevant packages, freeing them from the need to bid on *all* potentially relevant packages. Still, more bids may be needed than those that emerge during the clock stage. To accommodate this, the CCA adds a supplementary round, in which bidders can make many additional package bids. These are restricted, however, to be consistent – in terms of revealed preference – with some or all of the bids made during the clock rounds. The supplementary round completes the bid collection procedure.

[6]First-price auctions have been used in such contexts in countries including France and Norway, and Vickrey auctions have seen limited use in Canada.

Just as in first-price and Vickrey designs, the auctioneer then accepts as winning the set of bids – at most one per bidder – with the largest total bid value. The prices are most often computed using a "core-selecting" pricing rule that has, in practice, led to prices closely approximating Vickrey prices, given the submitted bids.

The CCA has been controversial for several reasons. One challenge has been that, just as in a Vickrey auction, different bidders may pay very different prices for identical packages. Another problem is that the auction game has been shown to have many equilibria, and in some of those an aggressive bidder makes licenses more expensive for its competitors and succeeds in buying more spectrum than passive bidders. More generally, because the Vickrey prices are set by losing bids, the CCA design creates opportunities for purely strategic behavior that is at odds with auction efficiency. Combinatorial designs also present other complications for bidder participation – in particular, the revealed preference rules used to limit the bids that can be made in the supplementary round are complex and difficult for bidders to manage.

5 Descending Clock Auctions

In the previous section, we described failures of the substitutes condition arising from bidder valuations; those failures destroyed some of the good properties of the Deferred Acceptance Algorithm, leading spectrum regulators to adopt entirely new designs. In a procurement auction context, the substitutes condition can also fail in another way, when the suppliers' offerings are not substitutes from the perspective of the auctioneer. This situation arose in the two-sided U.S. Broadcast Incentive Auction of 2016-2017 and led to the development of a new class of deferred acceptance algorithms: (generalized) descending clock auctions.

The economic problem to be solved arose from two market trends. One was that over-the-air (OTA) TV broadcasting was declining in economic value, as more consumers used cable, satellite and Internet-based services, rather than watching shows by radio broadcast. Another was that, since Apple's introduction of the iPhone, the use of spectrum for mobile broadband was increasing dramatically. Because of the intricate set of interference constraints governing each of these uses, the transition called for a coordinated change in which certain TV channels would be cleared from use for TV broadcast and made available for mobile broadband.

The market design needed to respect a constraint that each TV broadcaster had the right to continue to serve viewers in its broadcast area. Nearly every broadcaster, however, had different rights from any other. The viewers reached by each broadcast tower depended on its location, broadcast power and orientation, so the number of channels that would be required to allow a set of stations to continue OTA broadcasts depended on the precise combinations. In short, from the perspective of the auctioneer that hoped to clear a uniform band by buying out a set of broadcasters, the substitutes condition failed badly.

To promote value-enhancing trades, the Incentive Auction included several components: a voluntary *reverse auction* to identify a set of stations to acquire and prices to pay them; a *forward auction* to sell the mobile broadband licenses created with the vacated spectrum; a *reassignment plan* to assign channels to stations that would continue to broadcast on-air; and an overarching *market clearing procedure* to determine how many channels would be cleared.[7]

When the auction finished in 2017, winning broadcasters were paid about $10 billion to relinquish their broadcast rights and mobile operators paid the government about $20 billion for the mobile broadband licenses. This auction was a central pillar of the National Broadband Plan, ultimately freeing 84 MHz of spectrum to be used for broadband and other services throughout much of the United States.

The process used by the auction had to solve several novel problems. First, it was unclear in the period before the auction exactly what rights broadcasters had to the already-licensed spectrum. Licenses could not be simply revoked, but if every broadcaster had the right to continue broadcasting on its historical channel, then the need to clear a contiguous block of channels and the same channels in every city would lead to an intractable hold-out problem. The FCC and Congress ultimately decided that TV stations that chose not to sell their license rights could be reassigned to other channels, with the auction proceeds covering their retuning costs. With no TV station able to block the assignment of any particular channel, the hold-out problem was eliminated, and the stage was set for price competition among stations that wished to sell their rights within the same broadcast area.

Second, the auction process needed to determine how many channels to clear from TV broadcast uses. Each TV broadcast channel occupies a 6 MHz spectrum frequency. For any given number of channels, the government would try to buy spectrum rights sufficient to clear that many channels. Then, it would create and sell new mobile broadband licenses using the same frequencies. The revenues from the sale were required to cover the cost of acquiring licenses plus certain other expenses, including the cost of retuning the TV stations that continued on-air.

To decide how many channels to clear, the FCC announced opening prices before bidding started and required each station to declare whether it would participate in the reverse auction, which included a commitment to sell its rights at the opening price. The prices offered to each participant could only fall in the subsequent rounds of the reverse auction, so the opening price represented the most that a station could possibly receive to cease broadcasting. Using those participation decisions, the FCC determined the maximum number of channels that could possibly be cleared, and opened the auction using that quantity as both its demand in the reverse auction and the available supply in the forward

[7]Note that the term ‘clear’ is used here with two different meanings. *Market clearing*, used in the same sense as in previous sections, was satisfied for the overall Incentive Auction when the supply of spectrum sold in the reverse auction matched the demand for new spectrum licenses in the forward auction. Separately, the reverse auction *cleared* TV channels for new uses by paying winning stations to cease broadcasting.

auction. It then ran the two auctions. If the revenues from the forward auction were insufficient to cover the costs of the reverse auction plus the additional costs, then the number of channels to be cleared would be decreased and the process continued. Because the auction to buy broadcast rights used a certain descending clock auction and the sale of broadband licenses used an ascending clock auction, the prices in each auction were able to proceed monotonically as the clearing target was adjusted over time.

A third novelty of the Incentive Auction was the design of its reverse auction. The FCC staff had initially proposed adapting the Vickrey auction to this context, based on its status as the unique strategy-proof mechanism that always selects efficient outcomes. That advantage is important, but the Vickrey auction was nevertheless rejected in favor of a novel descending clock auction – an adaptation of a deferred acceptance algorithm – which was expected to perform much better for this auction problem.

5.1 Seven Weaknesses of the Vickrey Auction

Why not use a Vickrey auction, given its celebrated theoretical properties? Offsetting the Vickrey auction's advantages are seven disadvantages – each of which can be overcome by an alternative auction design.

Computational Feasibility. The first disadvantage is computational. If broadcast rights were purchased using a Vickrey auction, the first steps would be to collect bids and compute the set of stations with the lowest total bids, subject to the constraint that the set would be sufficient to clear the desired number of channels. For the Broadcast Incentive Auction, even checking just the *feasibility* of clearing a set of channels with certain station rights is an extremely challenging computational problem. To understand the reason, consider a simplification of the actual channel assignment problem.

Suppose that the TV interference constraints that need to be satisfied can be represented by a simple graph. Each node in the graph is a TV station that will continue to broadcast; each arc represents a pair of stations that cannot be assigned to the same channel without creating unacceptable interference. For the Broadcast Incentive Auction, before any clearing, there were about $2,400$ nodes (stations) and $137,000$ arcs (interference overlaps) in that graph, and the auction ultimately cleared about 200 stations. To check whether a particular set of stations can continue to broadcast using a given set of channels, it is necessary to find an assignment of nodes to channels such that no two connected nodes are assigned to the same channel, or to show that no such assignment exists. Replacing the word "channel" by "color," this is a version of the famously hard *graph-coloring problem.* The graph-coloring problem is *NP-complete*, which implies that, for any algorithm, and under standard assumptions, there is a sequence of such problems of increasing size such that the solution time grows exponentially in the problem size. In practice, this means that some large problems can be intractable: there may be no algorithm by which a fast computer can solve all of them, even with years of computation.

In the actual setting of the Incentive Auction, interference between two stations depended on the exact channels each is assigned, making the problem even more complex. There were also additional constraints to prevent any station near the U.S. border from interfering with Mexican and Canadian broadcasters. In total, there were about 2.7 million detailed constraints, which made the scale of the feasibility checking problem even larger. The FCC's computational experiments confirmed that this problem was sufficiently hard that a Vickrey auction, which is based on such a minimization of total bids, would not be viable.

Bidder Trust. The second reason that a Vickrey auction was inappropriate for this application is that it requires bidders to place too much trust in the auctioneer. Each bidder would effectively be told: "Tell us your station value. We will tell you whether your bid was winning and, if so, how much you will be paid. Unfortunately, we cannot guarantee that our computations will be accurate, because the computations are too difficult. Also, the data on which they are based are, by law, confidential, so you will be unable to check our computations." Meeting the goal of clearing a large number of channels required the participation of as many broadcasters as possible, so if many bidders had been deterred by concerns of trust, it would have substantially undermined the success of the auction.

Winner Privacy. A third concern is that the Vickrey auction does not "preserve winner privacy" – that is, all bidders in the auction would have been asked to report the lowest price that they would accept. In past second-price reverse auctions, commentators have often highlighted to the public that some winning bidders (those ultimately selected to sell an item or provide a service) had offered to accept a certain low price but were actually paid a much higher price. A similar objection would likely apply to every winning bidder, making managing public reception of the auction a significant challenge.

Budget Constraints. Fourth, the FCC had decided to clear as many channels as possible in the reverse auction, subject to a budget constraint determined by the revenues from the forward auction. This imposes a budget constraint on the reverse auction itself. If the reverse auction yields a cost for clearing a particular number of channels that exceeds the budget, then the auctioneer must reduce that number and continue the auction. The Vickrey auction is not consistent with any budget condition; it chooses its outcome based on maximizing an objective on some constraint set that depends on the allocation, but not on the Vickrey prices, which are determined in part by the values of certain losing bid combinations. If the FCC had tried to conduct a Vickrey auction and then to rerun it for a different channel target in case costs were too high, it would no longer have been a Vickrey auction – nor would it have possessed the Vickrey auction's desirable strategy-proofness and efficiency properties.

Group Strategy-Proof. A fifth disadvantages is that the Vickrey auction is not group strategy-proof. Consider an example with only three stations A, B and C, with station values v_A, v_B and v_C. Suppose that the coverage of stations A and C overlaps, as does the coverage of stations B and C. Before the auction, stations A and B broadcast simultaneously on channel #1, while station C uses channel #2.

If the auctioneer wants to clear one of the two channels for resale and leave the remaining OTA broadcaster(s) in the other channel, then it has two options: it can buy rights from both A and B to clear channel #1, which is efficient whenever $v_A + v_B < v_C$, or it can buy the rights of station C to clear channel #2. If A and B win, they are taken off the air and are paid the Vickrey prices $p_A = v_C - v_B$ and $p_B = v_C - v_A$, while station C receives $p_C = v_A + v_B$ if it wins. If stations A and B collude and both bid 0, they would win and each receive payment v_C. This is profitable for them whenever $\max\{v_A, v_B\} < v_C$, which can happen even if stations A and B would lose the auction under truthful bidding. Vickrey auctions are rare in their vulnerability to profitable collusion (even without compensating cash transfers) by losing bidders.

Price Competitiveness. Sixth, the Vickrey auction is not generally price competitive: there can be a collection of *losing* bidders sufficient to clear the channels who would be willing to accept a lower total price than the Vickrey prices. In the example above, if A and B win, then $p_A + p_B = 2v_C - v_A - v_B > v_C$, such that station C would have been willing to cease broadcasting at a price less than $p_A + p_B$. Vickrey prices are always competitive when the products are substitutes, so this may be regarded as a problem associated with complementarities.

Trading Off Efficiency and Cost. Seventh, the Vickrey auction is a single, fixed design that aims only to maximize efficiency. Although it is easy to extend the Vickrey design to include maximum prices for each station, cost minimizing designs require much more: they must replace the objective of maximizing total value in each realization by one of maximizing the sum of "virtual costs."

5.2 Avoiding the Seven Weaknesses

For the case in which each bidder has a single item to sell, the class of descending clock auctions avoids all seven weaknesses of the Vickrey auction. Like SMR auctions, descending clock auctions share some characteristics with the Deferred Acceptance Algorithm, but with additional variations. They allow auctions to be designed that overcome the seven weaknesses of the Vickrey auction design, as discussed above. The FCC design – used to procure broadcast rights from TV stations in the Incentive Auction – is but one member of this class.

Mechanism 5.1 (*Descending Clock Auction*)**.**

A descending clock auction evolves over a sequence of rounds, numbered by $t = 1, 2, \ldots$. Let N denote the set of initial bidders. At the beginning of each round t, there is some set of bidders $A(t) \subseteq N$ who are the *active bidders* that have not yet rejected any price offer and are eligible to become winners, while the other bidders $Z(t) \equiv A(t)^C$ have rejected a price offered before time t and have exited the auction. A *history* of the auction up to the start of round t is a collection $A^t = \{A(1), \ldots, A(t)\}$ where $A(t) \subseteq A(t-1) \subseteq \cdots \subseteq A(1) = N$, and H is the set of all possible histories.

Each descending clock auction is characterized by a pricing rule $p : H \to \mathbb{R}_+^N$, which maps each possible history A^t into a vector of prices for round t. In a *descending* auction, the prices in each round are no higher than in the previous round: $\forall (t \geq 2, A^t), p(A^t) \leq p(A^{t-1})$. If a bidder is active and its price changes, then it may choose either to *remain*, in which case it will again be active in the next round, or to *exit*, becoming inactive. An active bidder whose price does not change is offered no choice and also remains active. Each bidder observes only its own price, so its strategy in the auction is a mapping from a sequence of prices into a decision to *remain* or *exit*.

The auction is closed at the beginning of the first round T in which the specified prices do not change: $p(A^T) = p(A^{T-1})$. At this point, the winners of auction p are the bidders who are still active – that is, those in the set $A(T)$ – and they sell at the final prices $p(A^{T-1})$.

The quoted prices are relevant only to active bidders, so restricting p to allow price reductions only for active bidders does not affect the economic or game-theoretic analysis of the auction outcome. Any pricing function p satisfying all these restrictions is a descending clock auction, and any two different price functions p and p' satisfying the restrictions are different auctions, in the sense that for some bidder strategies, they lead to different outcomes.

For our game-theoretic analysis, we assume that each seller in the reverse auction plans to sell only a single item and has only one option: to sell or not. The designers of the Incentive Auction believed that this single-item description would apply to the vast majority of likely winners (successful sellers) in that auction, because the large station groups mainly included more valuable stations that were less likely to ultimately be sold.[8] For this single-item case, the class of descending clock auctions allows the auctioneer to avoid each of the seven weaknesses of the Vickrey auction.

Computational Feasibility. The computational challenges of any particular descending clock auction p arise from the complexity of computing its pricing function. For the Broadcast Incentive Auction, to ensure that the final outcomes would satisfy the non-interference constraints, the computations considered the stations one by one, as if each at time t the price were decreased for just one station. If the software determined that it would be feasible to "pack" the currently considered station – meaning to assign channels, without violating feasibility constraints, to a set of stations consisting of the current station and the other stations that either previously exited the auction or did not participate – then the price offered to that station would be reduced. Otherwise, the station would be considered infeasible to pack and its price would remain forever unchanged.

During the auction, each active station might be checked many times as its price was reduced over time; there were expected to be roughly $10,000$ feasibility checking operations required. To make the process fast enough for human

[8]The actual Incentive Auction also included other options for some station owners, allowing them to swap their UHF broadcast rights for VHF rights plus compensation, which are options that we do not discuss here.

bidders, the FCC limited the time allowed for feasibility checking by allotting a maximum time – typically about 60 seconds – to each problem. Since feasibility checking is NP-complete, the checking software would sometimes report that allotted time was not enough: it had run out of time and could not decide whether it was feasible to pack the station. Such a report was treated the same as if packing were proved to be infeasible: the station's price would not be reduced and would remain frozen for the remainder of the auction. With these specifications, the auction was guaranteed to finish all of its computations in a reasonable amount of time and to output a feasible set of stations to continue over-the-air broadcasting. Thus, *computationally feasibility* was guaranteed.

Theorem 5.2. *Every descending clock auction accommodates computation time limits. That is, for every descending clock auction p, there exists another descending clock auction p_C such that $p_C(A^t) = p(A^t)$ whenever that pricing rule can be computed within the time limit and such that if p always leads to feasible solutions after every history of feasible solutions, then p_C does so as well. A computationally feasible descending clock auction can be constructed by setting $p_C(A^t) = p(A^{t-1})$ when computation time expires.*

Thus, the descending clock auction design can tolerate small failures in computation. The Vickrey auction is not similarly robust: its desirable properties require the precise calculation of all winning prices.

One might nevertheless expect the NP-completeness of the packing problem to affect the quality of the final outcome. NP-completeness, however, is a worst-case criterion, meaning roughly that it is impossible to construct a feasibility checker that solves *all* of the desired problems quickly. However, the Incentive Auction software team was able to build a customized feasibility checker that solved about 99% of the problems in the typical allotted time of about 60 seconds per problem.

How did the 1% checking failures affect performance? When packing a station is feasible but the checker fails to discover that, then the auction price for that station could have been reduced but was not, such that the station remains active and may eventually become winning (by being paid to cease broadcasting). The consequence could be that the station is paid an unnecessarily high price, or that it is purchased to clear a channel when it need not have been purchased at all. Most feasibility checking failures, however, do not happen early in the auction when prices are high, because few stations have exited at that point and packing additional stations is easy. Instead, failures mostly happen late in the auction after many stations have exited, when packing becomes hard. Consequently, the expectation was that the 1% checking failures would have only a modest effect on efficiency and cost.

Bidder Trust. Bidders in any descending clock auction face a simple decision each round: whether to agree to sell their item at the stated price. If a bidder chooses not to sell, it exits the auction; if it chooses to continue bidding, its price can only decrease. No matter what the other rules may be or how others may bid, the bidder's dominant strategy is *honest bidding*, which dictates that it continue if its clock price exceeds its value and exit otherwise.

In the language of Li (2017), every descending clock auction is obviously strategy-proof (see chapter **??**). This means that at any information set reachable by honest bidding, the lowest payoff that the bidder can receive on any branch by continuing to bid honestly is at least as high as the highest payoff that it can get on any (possibly different) branch by some deviation.

Theorem 5.3. *Every descending clock auction is obviously strategy-proof. The obviously dominant strategy σ has the bidder remain if the price exceeds its value and exit if the price is less than its value.*

Proof. At every information set, following the strategy σ always leads to a non-negative payoff, while deviating leads to a non-positive payoff. □

Obvious strategy-proofness eliminates much of the bidder's need to trust the auctioneer. While no bidder can verify whether the auctioneer is following the announced rules correctly, it can easily follow whether its prices decrease from round to round – and this alone determines its optimal strategy. The bidder's strategy does not depend on whether the auctioneer has erred or cheated in undetectable ways.

Winner Privacy. Unlike a Vickrey auction, each winning bidder in a descending clock auction p reveals just enough information about its value during the auction to become a winner. All that an observer learns about a winning bidder from its bids in the auction is that the winner is willing to accept the final price, which is precisely the minimum information to determine that the bidder should be part of the final allocation. This condition can be described as (unconditional) *winner privacy.*

Theorem 5.4. *Every descending clock auction satisfies unconditional winner privacy.*[9]

Group Strategy-Proof.

Descending clock auctions are weakly group strategy-proof, which means that in the absence of side payments among participants, there exists no coalition of bidders who can all do strictly better for themselves by deviating to play any other joint strategy. The proof of this follows from obvious strategy-proofness. Consider the bidder who would first be called upon to deviate: the best outcome that bidder can achieve by deviating is no better than the worst it could get from continuing to bid honestly, so the deviation payoff cannot exceed the equilibrium payoff.[10]

Theorem 5.5. *Every descending clock auction is weakly group strategy-proof.*

Budget Constraints. Each descending clock auction that can achieve a desired allocation when a given budget constraint is not binding can be adapted

[9] Furthermore, if a strategy-proof mechanism satisfies unconditional winner privacy, then the same allocation can be implemented by a descending clock auction.

[10] By a similar argument, every obviously strategy-proof mechanism is group strategy-proof.

to a new descending clock auction that does the same and accommodates the budget constraint. Suppose we are given a descending clock auction p and a budget B. Define p_B to be a *budget-respecting extension* of p for budget B if p_B is a clock descending auction for which the total cost can never exceed B, and; for any value profile v for which p realizes total cost less than B, the allocations resulting from p_B and from p coincide.

The construction of one such p_B from p is easy and intuitive, so we can describe it in terms of the Broadcast Incentive Auction. In any round before the terminal round of p, p_B sets the same prices as p. If in the terminal round for p, the total prices exceed the budget B, then decrease all prices by \$1, and iterate until the total prices of active bidders are no more than B. Other budget-respecting extensions of p can be described similarly, but using different rules for price reductions after what would have been the terminal round for p. All such extensions are still descending clock auctions, and thus the other theoretical properties continue to hold. This contrasts with the Vickrey auction, for which no similar construction is possible.

Theorem 5.6. *Given any descending clock auction p and any budget $B > 0$, there is a descending clock auction p_B that is a budget-respecting extension of p.*

Price Competitiveness. A strategy-proof direct mechanism is *price competitive* if, for every value profile v, the set of winning bidders W and the prices they are paid coincide with those of some full-information Nash equilibrium of a first-price auction (with the same winner selection rule) in which every losing bidder bids its value. For any descending clock auction, there is a corresponding direct mechanism in which the bidders report their values and the mechanism plays the obviously dominant strategies on behalf of each bidder.

Theorem 5.7. *Every descending clock auction is price competitive.*

Proof. Unconditional winner privacy implies that the only information the algorithm uses about winning bidders' values is that they are no higher than the final prices, so if the winning bidders were able to report values equal to their prices instead of reporting truthfully, the algorithm would output the same set of winners.

In the first-price auction, the prices are equal to the bids. Winner privacy in the clock auction implies that if any winning bidder were to raise its bid, it would become a loser, so such a bid would be an inferior response to the strategies of the others. If a winning bidder were to reduce its bid, its price would fall, so that, too, is an inferior response. Hence, winning bidders' bids in the first-price auction are best responses to the other bidders' bid profile. By monotonicity, the losing bidders that increase their bids are still losers and losing bidders that reduce their bids cannot earn a positive payoff. So, the losing bidders, too, are playing best responses. □

Cost Minimization and Other Objectives. Next, we evaluate when and how descending clock auctions can be designed to achieve nearly maximum efficiency or minimum total procurement cost (or to approximately optimize other objectives) subject to incentive and participation constraints. Since all descending clock auctions are obviously strategy-proof and pay zero to losing bidders, the incentive and participation constraints are always satisfied.

In the Broadcast Incentive Auction, each TV station bidder becomes either a winner in set W or a loser in set L. After the auction, all the losers must be assigned to channels to continue OTA broadcasting, and the auction algorithm must ensure that such an assignment does not create unacceptable interference among those losers. The winners in the auction are paid to give up their broadcast rights.

Informally, this descending clock auction is an algorithm for a *packing problem*, in which losing bidders must be packed into a given set of broadcast channels. The steps of this algorithm are *greedy* – they select losing bidders irreversibly and without regard to the order of future exit decisions – and the full procedure is a *reverse* greedy algorithm because it iteratively selects *losers* rather than winners. In matching theory, algorithms that work by greedily rejecting individual offers are more commonly called *deferred acceptance algorithms.*

Formally, an instance of a packing problem $(v, s, \mathcal{F}, \omega)$ consists of a profile $v \in [0, \bar{v}]^N$ of station values, a profile $s \in \mathcal{S}^N$ of observed station characteristics (for example, s_n might encode station n's location, broadcast power, pre-auction channel, and population coverage, as well as the distribution from which v_n is drawn), a set $\mathcal{F} \subseteq \mathcal{P}(N)$ of feasible packings with the free disposal property that $Z \subset Z' \in \mathcal{F} \Rightarrow Z \in \mathcal{F}$, and a value $\omega(v_n, s_n)$ of each station that is successfully packed. The problem is to solve:

$$\max_{L \in \mathcal{F}} \sum_{n \in L} \omega(v_n, s_n)$$

The problem of maximizing the total value of the stations that will continue OTA broadcasting is a packing problem with $\omega(v_n, s_n) = v_n$. To minimize expected procurement costs, we adapt Myerson's optimal auction analysis to this problem. In outline, the first step is to determine the prior distribution of values for each station, which depends on its known characteristics s_n. The virtual value then depends on that distribution and on v_n, so it can be written as $\omega(v_n, s_n)$. Myerson showed that the expected total virtual cost of the winners is equal to the expected total payment to the winners. Consequently, the expected-cost minimizing auction maximizes the total virtual values of the losers: it solves the problem shown above.

Greedy algorithms can be applied to packing problems to approximate the optimal solution. To illustrate, consider the classic *knapsack problem*, in which each item has a value v_n and a size s_n. The collection of feasible sets consists of those for which the total size of the items does not exceed the capacity C of the knapsack: $\mathcal{F} = \{Z \subseteq \mathcal{P}(N) | \sum_{n \in Z} s_n \leq C\}$. Finding the optimal packing for the knapsack problem is NP-hard, but a greedy algorithm for this problem is fast and often performs well. This greedy algorithm first ranks the items according to $\frac{\omega(v_n)}{s_n}$ from largest to smallest and then adds the elements to the knapsack in that order, one at a time. If adding any particular element n would cause the knapsack constraint to be violated or would select an item with $\omega(v_n) < 0$, the algorithm skips that element and continues to the next. This procedure always leads to a feasible packing, with a total objective that is at least $1 - \frac{\max_n s_n}{C}$ of the optimum.

In the auction context, suppose that we adopt a direct mechanism and an algorithm that selects the losers Z using a reverse greedy algorithm. To do this, we specify a ranking function $r(v_n, s_n)$. Without loss of generality, we can limit the range of r to be $[0, 1]$. The algorithm is initialized by setting $Z = \emptyset$. It then applies the following steps iteratively:

1. Set $\hat{Z} = \{n \notin Z | Z \cup \{n\} \in \mathcal{F},\, r(v_n, s_n) > 0\}$.
2. Terminate if $\hat{Z} = \emptyset$.
3. Set $\hat{n} \in \arg\max_{n \in \hat{Z}} r(v_n, s_n)$ (breaking any ties in favor of the station with the lower index n) and update Z to $Z \cup \{\hat{n}\}$.

For strategy-proof implementation, the winner selection rule must be "monotonic," so the greedy ranking function must be non-decreasing in v_n. Clock auctions as we have defined them have discrete price decrements, but by using sufficiently small price increments, they can approximate any monotonic ranking function.

Theorem 5.8. *For every reverse greedy algorithm associated with a ranking function $r(v_n, s_n)$ that is non-decreasing in its first argument, there is a family $\{p^k\}$ of k-round descending clock auctions, such that for every value profile v and all k sufficiently large, straightforward bidders exit auction p^k in the same order as the selection order of the reverse greedy algorithm.*

Proof. The family of descending clock auctions is constructed as follows. For auction p^k, time is indexed by $t \in \{0, \frac{1}{k}, \frac{2}{k}, ..., \frac{k-1}{k}, 1\}$. The set of selected items before time t is denoted by $Z^k(t)$ and the price vector then by $p^k(t)$. The auction initializes with $Z^k(0) = \emptyset$ and all station prices given by the high price $p^k_n(0) = \bar{v}$. Prices are set as follows:

$$p^k_n(t) = \begin{cases} q_n(t) \equiv \sup\{q \in [0, \bar{v}] | r(q, s_n) \leq 1 - t\} & \text{if } Z^k(t) \cup \{n\} \in \mathcal{F} \\ p^k_n(t - \frac{1}{k}) & \text{otherwise} \end{cases}$$

Thus, the price offered to a station is reduced only if the station can be packed; and if reduced, the price is set to control for the effect of the station's characteristics s_n on its ranking relative to other stations. (To break ties, each round consists of N sub-rounds, in which the clock prices of stations $1, ..., N$ are adjusted sequentially, beginning with station $n = 1$. If any station n exits in a sub-round, then $Z^k(t)$ is updated, the prices $p^k_m(t)$ for the stations $m > n$ are recomputed, and the processing continues with station $n + 1$.) By inspection, this pricing rule defines a descending clock auction p^k.

A straightforward bidder exits at time t if its price is then less than its value, so if station n exits at time t, then $v_n > p_n^k(t) = q_n(t)$. If there are no ties in the auction and station m can be feasibly packed but has not yet exited at t, then $v_m \leq p_m^k(t) = q_m(t)$. Hence, $r(v_n, s_n) \geq r(q_n(t), s_n) > r(p_m^k(t), s_m) \geq r(v_m, s_m)$. Thus, without ties, the clock auction packs feasible stations in the same order as the reverse greedy algorithm. If there are no ties in the reverse greedy algorithm, then for k sufficiently large, there are no ties in the clock auction. If there are ties in the greedy algorithm, then they are broken in the same way in the clock auction and, inductively, the same applies to each subsequent choice. □

According to this theorem, every reverse greedy algorithm can be approximated by a descending clock auction that leads to exactly the same packing with arbitrarily high probability, but the full class of algorithms that can be implemented by clock auctions is much larger. As we have seen, the set of stations that is packed by a descending clock auction can depend on the prices that are paid to bidders, which is helpful for meeting budget constraints, and clock auctions remain computable even when the exact feasibility checking required by the greedy algorithm is impracticably difficult.

For minimizing expected total procurement costs, we adapt the famous formulation of Myerson's Optimal Auction Design. Assume that the bidders' values v_n are independently distributed with known distributions F_n with strictly positive densities f_n on an interval $[0, \bar{v}]$. Adapting Myerson's analysis to a procurement problem, define the *virtual cost* for a bidder with value v_n to be $C_n(v_n) = v_n + \frac{F_n(v_n)}{f_n(v_n)}$. Myerson's lemma, rewritten for the procurement case, states that for any Bayesian incentive-compatible mechanism with zero payments to losing bidders, the expected total payment to winning bidders in W is $\mathbb{E}\left[\sum_{n \in W} C_n(v_n)\right]$. This lemma applies to all descending clock auctions. For brevity, we limit attention to the case in which each $C_n(\cdot)$ function is increasing and continuous. For this case, as Myerson showed, the optimal mechanism selects winners for each realization of the value profile v according to:

$$x(v) \in \underset{\{W | W^C \in \mathcal{F}\}}{\arg\min} \sum_{n \in W} C_n(v_n)$$

Since $L = W^C$, an equivalent characterization of the allocation rule is:

$$x(v)^C \in \underset{L \in \mathcal{F}}{\arg\max} \sum_{n \in L} C_n(v_n)$$

By Theorem 5.8, every greedy algorithm for this latter problem can be approximated arbitrarily closely by a descending clock auction.

In the actual Broadcast Incentive Auction, the station value distributions F_n were not known, the assumption of independent station values was tenuous, and the interference constraints were not knapsack constraints, but the two preceding examples nevertheless guided the choice of an auction pricing function. The observable station characteristics were taken to be $s_n = (Pops_n, Links_n)$, where $Pops_n$ is the population reached by the broadcast signal and $Links_n$ is the number of station's links in the broadcast interference graph. By analogy to the knapsack problem, $Links_n$ was expected to be a size-like variable, measuring the difficulty of packing a station on-air. For the Myerson-like cost-minimization, $Pops_n$ parameterized the station-value distribution. The price for going off-air in any round t offered to any still-feasible station n was:

$$p_n^{FCC}(t) = Links_n^{0.5} Pops_n^{0.5} p_0^{FCC}(t),$$

where $p_0^{FCC}(t)$, called the "base clock price," would fall during the auction following a predetermined schedule. This particular pricing function was explained to participants by noting that Links and Pops were weighted equally, and that multiplying both Links and Pops by the same constant multiplies the price by the same factor. The auction design team predicted that the $Links_n$-adjustment would improve the quality of packing, just as the size adjustment does in the greedy algorithm for the knapsack problem. Similarly, the $Pops_n$-adjustment would continue to offer a higher per-Pop price to smaller stations, encouraging the smaller stations to be the ones to sell. Analysis of the Myerson-optimal auction with varied assumptions about the distributions of values suggested that this would be a robust characteristic of a cost-reducing auction.

6 Future Auctions

Spectrum allocation problems occupy an intermediate space between commodities markets and discrete matching markets and often incorporate elements of both. Prices are important to allocate spectrum quantities efficiently, but the licenses are sufficiently specialized that matching the right combinations of frequencies to users is also critically important. Well-designed spectrum auctions must anticipate challenges arising from thin markets and bidder incentives, as well as broader competitive issues and meaningful value complementarities. Market designers have addressed many of these challenges and must continue to do so as new technologies increase demand for spectrum use. This will require innovative, efficient allocations that can also accommodate important regulatory concerns.

The theory that has developed to support the design of spectrum auctions may also find new applications in markets with similar issues of heterogeneous matching with prices and complementary goods, including markets for physical and natural resources. Almost all physical resources are heterogeneous in quality, time and place. Like spectrum, the value and availability of water, minerals and energy depend on all those factors. In a localized context, electricity may seem to be fully fungible, and users and producers can be connected using a common commodity market. From the perspective of a grid regulator, however, electricity in the morning in Los Angeles and electricity at night in San Francisco are not substitutes; more reliable suppliers call for greater precautions than less reliable ones; users may prefer green energy sources to others; and so on.

In many markets including those for spectrum licenses, a crucial step to improving allocative efficiency is to identify the appropriate division of the resource – the *product* to be sold or matched. Careful product definition is important wherever goods can be standardized to varying degrees. In some cases, making finer distinctions between products can improve efficiency; some online advertising impressions were once treated as interchangeable commodities, but today most high-value impressions are uniquely auctioned and matched. In others, including spectrum allocation, fine-grained differentiation may be theoretically appealing but present meaningful complications to bidders that make reduce competition for each product and make markets thinner. Ongoing work on new bidding languages might allow bidders to compactly represent complex preferences, thickening the overall market. While it is possible to add flexibility in ascending auctions to bid for licenses that are not exact substitutes, the existing adaptations are imperfect: there is no way to bid to match exactly a competitive demand for licenses that are complements.[11] The next generation of auction and matching algorithms may introduce new languages that can accommodate novel and more complex resource allocation challenges.

7 Bibliography

Much of the theory initially used to analyze spectrum auctions is the same as matching theory, including the seminal results of Gale and Shapley (1962) on one-to-one matching without prices and Shapley and Shubik (1971) on one-to-one matching with prices, the unifying results of Kelso and Crawford (1982) within the many-to-one firm-worker model, and the extension to contracts of Hatfield and Milgrom (2005) and Fleiner (2003). The algorithmic properties of auctions for the case of substitutes were further developed by Demange et al. (1986), Milgrom (2000), and Gul and Stacchetti (2000). These papers, however, paid little attention to the incentive issues for dynamic auction mechanisms and the special problems associated with pervasive failures of the substitutes condition.

Section 3 emphasizes broader market issues as collected by Klemperer (2002). Riedel and Wolfstetter (2006) provided a game-theoretic analysis of the ascending clock auction, highlighting a vulnerability to collusive-seeming outcomes that has been important in practice. The analysis of Treasury auctions and incentive challenges in uniform-price sealed-bid auctions is taken from Back and Zender (1993).

Bulow et al. (2009) offered an expanded account of FCC Auction 66, as summarized in Section 4. Ausubel and Milgrom (2006) provided results about the virtues and weaknesses of Vickrey auctions, while Bernheim and Whinston (1986) offered analysis of first-price auctions. The CCA design was first

[11] [PRM] The original SMRA design allowed bidders to withdraw some bids during the auction, but these withdrawals came with penalties or obligations for the bidder to cover at least part of the auctioneer's loss in case the withdrawn bids later result in a price reduction for those products.

proposed by Ausubel et al. (2006), and Day and Milgrom (2008) provided the core-selecting pricing rule. Levin and Skrzypacz (2016) analyzed the historical performance of the CCA, identifying problematic equilibrium solutions.

Despite what appears to be a vast difference between the mathematical structure of matching theory and the problems of the Broadcast Incentive Auction, the auction design can be understood as a kind of Deferred Acceptance Algorithm, in which bidders make a sequence of offers that are rejected and replaced until, finally, the bids that were never rejected become the winning ones. The auction design itself could not succeed without a careful design of spectrum rights, developed by Evan Kwerel and John Williams and codified by legislation, and new algorithms for NP-hard packing problems developed by a team headed by Kevin Leyton-Brown. Milgrom and Segal (2014) and Milgrom and Segal (2020) formalize and analyze the class of clock auctions from which the FCC's reverse auction was derived. Variations of their results form the basis of Section 5. Newman et al. (2020) reports computational experiments examining the realized efficiency of the Incentive Auction. Myerson (1981) provided seminal results in the theory of optimal auctions, and Li (2017) introduced the concept of obviously strategy-proof mechanisms.

References

[1] Ausubel, Lawrence M, and Milgrom, Paul. 2006. The lovely but lonely Vickrey auction. *Combinatorial auctions*, **17**, 22–26.

[2] Ausubel, Lawrence M, Cramton, Peter, and Milgrom, Paul. 2006. The clock-proxy auction: A practical combinatorial auction design. *Handbook of Spectrum Auction Design*, 120–140.

[3] Back, Kerry, and Zender, Jaime F. 1993. Auctions of divisible goods: On the rationale for the treasury experiment. *The Review of Financial Studies*, **6**(4), 733–764.

[4] Bernheim, B Douglas, and Whinston, Michael D. 1986. Menu auctions, resource allocation, and economic influence. *The quarterly journal of economics*, **101**(1), 1–31.

[5] Bulow, Jeremy, Levin, Jonathan, and Milgrom, Paul. 2009. *Winning play in spectrum auctions.* Tech. rept. National Bureau of Economic Research.

[6] Day, Robert, and Milgrom, Paul. 2008. Core-selecting package auctions. *international Journal of game Theory*, **36**(3-4), 393–407.

[7] Demange, Gabrielle, Gale, David, and Sotomayor, Marilda. 1986. Multi-Item Auctions. *Journal of Political Economy*, **94**(4), 863–872.

[8] Fleiner, Tamás. 2003. A fixed-point approach to stable matchings and some applications. *Mathematics of Operations Research*, **28**(1), 103–126.

[9] Gale, D., and Shapley, L. S. 1962. College Admissions and the Stability of Marriage. *The American Mathematical Monthly*, **69**(1), 9.

[10] Gul, Faruk, and Stacchetti, Ennio. 2000. The English Auction with Differentiated Commodities. *Journal of Economic Theory*, **92**(1), 66–95.

[11] Hatfield, John William, and Milgrom, Paul R. 2005. Matching with Contracts. *American Economic Review*, **95**(4), 913–935.

[12] Kelso, Alexander S., and Crawford, Vincent P. 1982. Job Matching, Coalition Formation, and Gross Substitutes. *Econometrica*, **50**(6), 1483.

[13] Klemperer, Paul. 2002. What really matters in auction design. *Journal of economic perspectives*, **16**(1), 169–189.

[14] Levin, Jonathan, and Skrzypacz, Andrzej. 2016. Properties of the combinatorial clock auction. *American Economic Review*, **106**(9), 2528–51.

[15] Li, Shengwu. 2017. Obviously strategy-proof mechanisms. *American Economic Review*, **107**(11), 3257–87.

[16] Milgrom, Paul. 2000. Putting Auction Theory to Work: The Simultaneous Ascending Auction. *Journal of Political Economy*, **108**(2), 245–272.

[17] Milgrom, Paul, and Segal, Ilya. 2014. Deferred-acceptance auctions and radio spectrum reallocation. Pages 185–186 of: *Proceedings of the fifteenth ACM conference on Economics and computation*.

[18] Milgrom, Paul, and Segal, Ilya. 2020. Clock auctions and radio spectrum reallocation. *Journal of Political Economy*, **128**(1), 1–31.

[19] Myerson, Roger B. 1981. Optimal auction design. *Mathematics of operations research*, **6**(1), 58–73.

[20] Newman, Neil, Leyton-Brown, Kevin, Milgrom, Paul, and Segal, Ilya. 2020. Incentive Auction Design Alternatives: A Simulation Study. Pages 603–604 of: *Proceedings of the 21st ACM Conference on Economics and Computation*.

[21] Riedel, Frank, and Wolfstetter, Elmar. 2006. Immediate demand reduction in simultaneous ascending-bid auctions: a uniqueness result. *Economic Theory*, **29**(3), 721–726.

[22] Shapley, L. S., and Shubik, M. 1971. The assignment game I: The core. *International Journal of Game Theory*, **1**(1), 111–130.