Spectrum Auction Design

Peter Cramton*

22 April 2009

Preliminary Draft

Abstract

Spectrum auctions are used by governments to assign and price licenses for wireless communications. The standard approach is the simultaneous ascending auction, in which many related lots are auctioned simultaneously in a sequence of rounds. I analyze the strengths and weaknesses of the approach with examples from US spectrum auctions. I then present a variation, the package clock auction, adopted by the UK, which addresses many of the problems of the simultaneous ascending auction while building on its strengths. The package clock auction is a simple dynamic auction in which bidders bid on packages of lots. Most importantly, the auction allows alternative technologies that require the spectrum to be organized in different ways to compete in a technology-neutral auction. In addition, the pricing rule and information policy are carefully tailored to mitigate gaming behavior. An activity rule based on revealed preference promotes price discovery throughout the clock stage of the auction. Truthful bidding is encouraged, which simplifies bidding and improves efficiency. Experimental tests and early auctions confirm the advantages of the approach.

(JEL D44, C78, L96. Keywords: auctions, spectrum auctions, market design, package auction, clock auction, combinatorial auction.)

1 Introduction

Spectrum auctions have been used by governments to assign and price spectrum for the last fifteen years. Over those years, the simultaneous ascending auction, first introduced in the US in 1994, has been the predominant method of auctioning spectrum. The auctions have proved far superior to the prior methods of beauty contests and lotteries (Cramton 1997, Milgrom 2004).

Despite the generally positive experience with auctions, several design issues have surfaced. Some were addressed with minor rule changes. For example, bidders' use of trailing digits to signal other bidders and support tacit collusion was eliminated by limiting bids to integer multiples of the minimum increment (Cramton and Schwartz 2002). However, many other design problems remain. In this paper, I identify these problems, and describe a new approach, the package clock auction, based primarily on the clock-proxy auction (Ausubel et al. 2006), which addresses the main limitations of the simultaneous ascending auction.

My focus here is on spectrum auction design, rather than spectrum policy more generally. Certainly communications regulators face many other critical challenges, such as how best to free up new spectrum for

* Department of Economics, University of Maryland. I thank my collaborators, Larry Ausubel, Robert Day, and Paul Milgrom for helpful discussions, as well as Nathaniel Higgins, Evan Kwerel, Thayer Morrill, Peter Pitsch, and Andrew Stocking. I thank the staff at Ofcom, especially Graham Louth, Director of Spectrum Markets, whose leadership and intellectual contribution were essential to the successful implementation of the package clock auction. I am grateful to the National Science Foundation and the Rockefeller Foundation for funding.

auction (Cramton et al. 1998), or whether an auction is needed at all (FCC 2002). For some allocations, it is better to set aside the spectrum for common property use. In particular, for applications that do not create additional scarcity, the commons model is better than the auction model. There are many examples of this: garage door openers, car locks, and other device controllers. The application requires little bandwidth or power, and thus, does not make the spectrum scarce. Scarcity problems are mitigated by operator separation. In contrast, mobile phones require much greater power and bandwidth, creating spectrum scarcity, and hence an auction is needed to assign the scarce resource among the competing operators.

Spectrum auctions to date have been long-term auctions in which the winner is granted a license of from 10 years to 25 years, with a strong expectation of renewal following expiration. One might think instead that a spot market for spectrum, much like a spot market for electricity, would be a more flexible and efficient instrument. Someday that will be true. But today's hardware, especially the handset, is not sufficiently flexible to accommodate a real-time spot market. Moreover, operators must make large specific investments in their networks. These long-term investments are better supported with a long-term license for spectrum, a critical input. Over the next twenty years increasingly flexible hardware will be introduced. Eventually it will make sense to organize the spectrum market much like the electricity market. The basic element will be a real-time spot market that establishes the price of bandwidth at a particular time and location. But for now, long-term spectrum auctions are both necessary and desirable.

One of the greatest challenges for the regulator is keeping up with the rapid technological development of wireless communications. Indeed, one of the main reasons for switching from beauty contests, to lotteries, to auctions was that beauty contests and lotteries were too slow. Wireless communications plays an essential role in modern economies, both in developed and developing countries. Slowing the pace of wireless innovation and development has large costs to economic growth. For this reason, regulators must do whatever they can to promote a competitive wireless industry. Allocating sufficient spectrum in a timely manner is paramount.

The package clock auction described here helps facilitate the spectrum allocation process by enabling the auction to determine how the spectrum is organized, what is called the band plan. Prior methods required that the regulator determine a fixed band plan before the auction began. As a result, before each auction there is a long regulatory process, much like the beauty contests of before, but with the companies lobbying for particular band plans, rather than for direct spectrum awards. This is the most time consuming and error prone element of the spectrum management process. Thus, the new approach promises to not only improve spectrum assignments, but to improve the band plans within which the assignments fit, and to do so with less delay.

From an auction theory viewpoint, spectrum auctions are both challenging and interesting. The government is auctioning many items that are heterogeneous but similar. Often there are competing technologies as well as companies to provide a wide range of communication services. As a result, the setting has a complex structure of substitutes and complements. This is among the most difficult auction settings seen in practice.

The goal for the government should be efficiency, not revenue maximization. The government should focus on ensuring that those who can put the spectrum to its highest use get it. Focusing simply on revenue maximization is short-sighted. Many steps such as technical and service flexibility, and license aggregation and disaggregation improve efficiency, and thereby improve revenues. But short-run revenue maximization by creating monopolies, which would create the highest profits before spectrum fees, and therefore would sustain the largest fees should be resisted. Indeed, competition, which ultimately will lead to greater innovation and better and cheaper services, will likely generate *greater* government revenues from a long-run perspective. The government can best accomplish this objective with an efficient auction, putting the spectrum to its best use.

The regulator may find it necessary to introduce spectrum caps or other preferences favoring new entrants so as to level the playing field between incumbents and new entrants. Incumbents include in their private value, the benefit of foreclosing competition, thus driving a wedge between social value and private value. In theory the regulator can correct this externality by favoring the new entrant, but in practice this has proven to be difficult. The FCC's experience with preferences for certain bidders—set asides, bidding credits, and installment payments—has been largely a failure at least with respect to mobile broadband communication, which is where most of the value lies.

In contrast, a good example of successful intervention was Canada's use of a set aside in its 2008 Advanced Wireless Services or AWS auction. As a result, multiple deep-pocketed new entrants came to the auction, and bid up the price of not only the set-aside block, but the non-set-aside blocks. The result was a much more competitive auction (with much higher revenues) and the introduction of some potentially strong new service providers going forward. The approach effectively broke up regional market-splitting by the dominant incumbents. Another successful intervention was the FCC's use of a spectrum cap in early broadband PCS auctions. The cap limited the quantity of spectrum any one operator could hold in a geographic area, addressing the potential market failure of limited competition in the market for wireless services.

Despite these successes in Canada and the US, the FCC's long and troubled history with bidder preferences is an important case study to other countries considering preferences for various parties. Installment payments proved especially problematic, as it led to speculative bidding, bankruptcy, and lengthy delay in the use of the spectrum.

In addition, the regulator must resist the temptation to force more "winners" than the market can efficiently support. Sometimes regulators fragment the spectrum and prohibit aggregation in the auction in an effort to create as many winners as possible. The upcoming 3G spectrum auction in India may be one example. Aggregation up to a suitable competitive constraint is preferred.

1.1 Three main points

There are three main points I wish to emphasize.

Enhance substitution. First in terms of the auction design, it is important to enhance the substitution across the items that are being sold. Enhanced substitution is accomplished through both the product design—what you are auctioning—and of the auction format. Often in the spectrum setting, the product design can be just as important as the auction design.

Encourage price discovery. Second, encouraging price discovery is extremely important. We need a dynamic process here, because unlike some situations, in the case of spectrum auctions, there is a lot of uncertainty about what things are worth. The bidders need to do a lot of homework to develop a crude valuation model, and they need the benefit of some collective market insights, which can be revealed in a dynamic auction process, in order to improve their decision-making. The nice thing about a dynamic auction is that through this price process the bidders gradually have their sights focused on the most relevant part of the price space. Focusing bidder decisions on what is relevant is in my mind the biggest source of benefit from the dynamic process. This benefit is generally ignored by economists, because economists assume that the bidders fully understand their valuation models, when in practice bidders almost never have a completely specified valuation model. Yes, they do a lot of homework, but there is still a lot of uncertainty about what exactly things are worth, and how they should be valuing the spectrum. The experience of the 3G spectrum auctions in Europe is a good example. The bids were based more on stock prices in a bubble situation, rather than on solid homework about values.

Induce truthful bidding. The third feature I wish to emphasize is the importance of inducing truthful bidding. This is accomplished in the auction design through an effective pricing rule and an activity rule. The two rules work together to encourage bidders to truthfully express preferences throughout the entire auction.

This truthful expression of preferences is what leads to excellent price discovery and ultimately an efficient auction outcome.

A variety of different pricing rules are used in practice. The two most common are pay-as-bid pricing, where you pay what you bid if you are the winner, and for a homogenous product, uniform pricing, where you pay the market clearing price. In the particular applications I am discussing here, we generally do not have clearing prices, because of strong complementarities and heterogeneous items. As a result, we need a new kind of pricing rule. The pricing rule that I will describe in detail later is a generalization of Vickrey´s second-price rule.

I now give a brief overview of the package clock auction. The approach may appear complex. Some amount of complexity is required given the complex economic problem at hand. Simpler versions, such as a simultaneous clock auction are possible in settings where all bidders intend to use the same technology. This may well be the case in developing countries that are conducting spectrum auctions for a particular use after the technology battles have been resolved from the experience in developed countries.

1.2 An overview of the package clock auction

The package clock auction is especially useful in situations where the regulator does not know which technology will make best use of the spectrum. In such cases, the auction itself can determine the ultimate band plan specifying how the spectrum is organized. Such an auction is said to be technology neutral, since it allows the competing technologies to determine the winning technologies, as well as operators. A good example is an auction that accommodates both paired and unpaired technologies, such as LTE and WiMAX, respectively. A package auction is essential in this case, since the two uses require that the spectrum be organized in fundamentally different ways. The package clock auction is an especially simple, yet powerful, auction that lets competitive bids determine the ultimate band plan.

The package clock auction has features to address each of my three main points.

First, the product design simplifies the products whenever possible. For example, if bidders primarily care about the quantity of spectrum won in a geographic area, we auction generic spectrum and the bidders bid for a quantity of spectrum in each area. This simplifies the auction, enhances substitution, and improves competition. The specific assignment of lots is determined in the last stage of the auction, once the critical decisions have been made (who won how much in each area). This approach also allows a technology neutral auction, which lets the spectrum be organized in different ways for the different technologies. Each bidder indicates the quantity of spectrum and the type of use in its bids. In this case, the first stage of the auction determines not only who won how much in each area, but the overall quantity of spectrum allocated for a particular use in the area.

Second, to encourage price discovery, the auction begins with a clock stage. Prices ascend for each product with excess demand until there is no excess demand for any product. This simple and familiar price discovery process works extremely well when bidders have incentives for truthful bidding. In the important case of substitutes, the clock stage determines an efficient assignment together with supporting competitive equilibrium prices. Moreover, complements are handled with no increase in the complexity of the clock process. Each bid in the clock stage is a package bid, so bidders can bid without fear of winning only some of what they need.

The bidders may find that they are unable to express preferences for all desirable packages in the clock stage, so following the clock stage is a supplementary round. Bidders can increase their bids on packages bid in the clock stage and submit new bids on other packages. All the clock stage bids and the supplementary round bids then are run through an optimizer to determine the value maximizing assignment of the spectrum. This is the generic assignment.

Third, to induce truthful bidding, the auction uses closest-to-Vickrey core pricing. The efficient assignment is priced to minimize the bidders' total payments subject to competitive constraints (no group of bidders has offered the seller more). In practice, this often implies Vickrey pricing, ensuring truthful bidding. However, because of complements, there may be one or more competitive constraints that causes the payments to be greater than Vickrey payments for some bidders. In this event, the smallest deviations from Vickrey prices are used. This maximizes incentives for truthful bidding subject to competitive constraints.

To induce truthful bidding throughout the clock stage, an activity rule based on revealed preference is used. This rule encourages bidders to bid in a manner consistent with profit maximization. Deviations from bidding on the most profitable package throughout the clock stage may impose a constraint on subsequent bids, either later in the clock stage or in the supplementary round. To simplify the auction, a simplified revealed preference rule is desirable that maintains the same one-dimensional structure of the traditional activity rule. In particular, every bid in the clock stage or the supplementary stage is constrained by at most a single constraint. This simplified rule encourages straightforward bidding—always bidding on the most profitable package—without complicating the auction.

Once the generic assignments are determined and priced, the specific assignment stage is run. Each winner submits top-up bids for each specific assignment that is better than the winner's worst specific assignment. The bids indicate the incremental value for each feasible alternative. Then an optimization program is run to determine the efficient specific assignment. Again the prices for the specific assignments are closest-to-Vickrey core prices. This concludes the auction.

This paper builds on well the developed literatures on auction theory and practice, especially combinatorial auctions and spectrum auctions. Much of the literature on combinatorial auctions is summarized in Cramton et al. (2006). The work of Ausubel et al. (2006), Ausubel and Milgrom (2006a,b), Day and Raghavan (2007), Day and Milgrom (2008), Day and Cramton (2008), Milgrom (2007, 2009), Parkes (2006), and Porter et al. (2003) is especially relevant. On spectrum auctions see Coase (1959) for the original proposal, Ausubel et al. (1997) on synergies, McMillan (1994), Cramton (1995, 1997, 2006), Klemperer (2004), and Milgrom (2004) on the performance of the simultaneous ascending auctions, and Brusco and Lopomo (2002) and Cramton and Schwartz (2002) on collusion.

I begin by describing some of the problems of the simultaneous ascending auction. Then I present the package clock auction, which retains the benefits, while addressing the weaknesses, of the simultaneous ascending auction. I emphasize two essential elements of the package clock auction: the pricing rule and the activity rule. Along the way, I summarize both experimental and field results with the package clock auction.

2 Simultaneous ascending auction

The workhorse for spectrum auctions since 1994 has been the simultaneous ascending auction, a simple generalization of the English auction to multiple items in which all items are auctioned simultaneously. Thus, unlike Sotheby's or Christie's auctions in which the items are auctioned in sequence, here all the items are auctioning simultaneously.

The process is as follows. Each item or lot has a price associated with it. Over a sequence of rounds, we ask bidders to raise the bid on any of the lots they find attractive, and we identify the standing high bidder for each lot at the end of every round. We continue this process until nobody is willing to bid any higher. This process was originally proposed by Preston McAfee, Paul Milgrom, and Robert Wilson for the FCC spectrum auctions. Since its introduction in July 1994, the design has undergone numerous enhancements, but the basic design has remained intact in its application worldwide for the vast majority of spectrum auctions.

An important element of the basic design is an activity rule to address the problem of bid sniping—waiting until the last minute to bid seriously. The rule adopted by the FCC and used in all simultaneous ascending auctions to date is a quantity-based rule. In short, the rule states, "if you want to be a big bidder at the end of the auction, you must be a big bidder throughout the auction." You must maintain a level of eligibility, based on the quantity of spectrum you are bidding for, in order to continue with that level of eligibility later on. Thus, you cannot play a snake-in-the-grass strategy where you hold back and wait, and then pounce late in the auction and win without making your true intent known until the last instant.

As mentioned, the simultaneous ascending auction has been with us for a long time. The FCC has conducted 72 simultaneous ascending auctions, since it was introduced in July of 1994. The FCC has gotten very good at conducting the auctions, and the design has worked reasonably well. Nonetheless, it is perhaps surprising how quickly inertia set in. The FCC was initially highly innovative in its choice of design the first time out of the block, but since then they have just made minor incremental improvements in response to obvious and sometimes severe problems with the original simultaneous ascending auction design.

Why has the design held up so well? It is an effective and simple price discovery process. It allows arbitrage across substitutes. It lets bidders piece together desirable packages of items. And, because of the dynamic process, it reduces the winner's curse by revealing common value information during the auction.

But the design does, and has been observed to have, many weaknesses.

- As a result of the pricing rule, there is a strong incentive for large bidders to engage in demand reduction—reducing the quantity demanded before the bidder's marginal value is reached in order to win at lower prices.
- Especially if there is weak competition, bidders have an incentive to engage in tacit collusion. The bidders employ various signaling strategies where they attempt to work out deals through the language of the bids. The goal of the strategies is to divvy-up the items among the bidders at low prices.
- As a result of the activity rule, there are parking strategies. A bidder maintains eligibility by parking its eligibility in particular spots that the bidder is not interested in and then moves to its true interest later.
- The simultaneous ascending auction is typically done without package bids. The bidders are bidding on individual lots and there is the possibility that a bidder will win some of the lots that it needs for its business plan, but not all. This exposure to winning less than what the bidder needs has adverse consequences on efficiency. Essentially the bidder has to guess. Either the bidder goes for it or not. When there are complementarities, this is a tough decision for the bidder to make. The bidder may make the wrong decision and win something it actually does not want.
- The lack of package bids also makes the simultaneous ascending auction vulnerable to hold up, which is basically a speculator stepping in and taking advantage of a bidder (Pagnozzi 2007). The speculator can make it clear to large bidders that it would be expensive to push him out of the way. As a result, the large bidders let the speculator win some desirable lots at low prices, and then the speculator turns around and sell them to the big players after the auction is over, and make some quick money. That is the holdup strategy. It is easy to do and effective. Preventing resale would reduce this problem, but resale is desirable in a rapidly changing dynamic industry.
- There is limited substitution across licenses, which is something I am going to emphasize. You might think that it would be easy to arbitrage across the lots, but in fact that is not the case. Especially in a large country like the United States, where the FCC splits up the frequency bands in different ways, geographically, and you can only bid on individual lots, rather than packages.

As a result of all these factors, the bidding strategies are quite complicated, which is nice if you happen to be advising bidders on spectrum auctions, but is a problem if you are a bidder, because you have to learn how to engage in all this complex bidding.

2.1 The US AWS and 700 MHz auctions

The difficulties in arbitraging across substitutes are best illustrated in the two most recent major auctions in the United States: AWS and 700 MHz.

Figure 1. The US AWS band plan: something for everyone

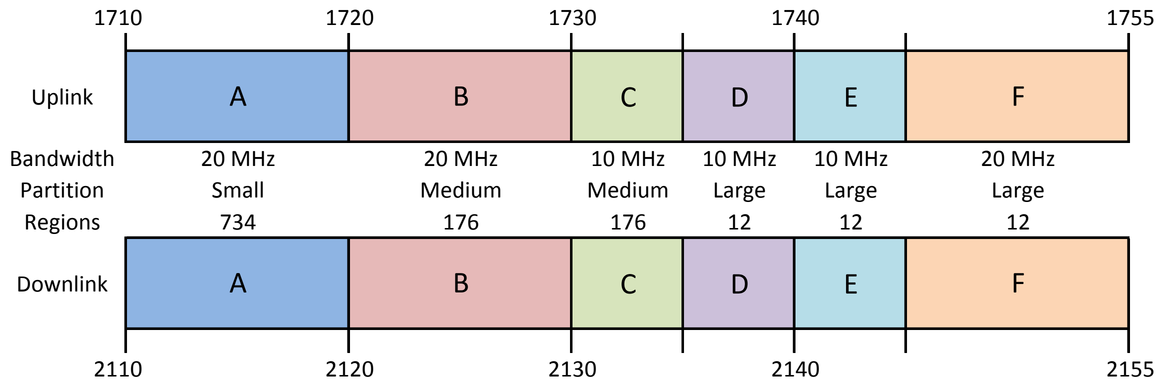

First the AWS, Advances Wireless Services, auction sold 90 MHz of spectrum in 161 rounds, and raised 14 billion dollars. As in all of its auctions, the FCC began the process by settling on a specific band plan (the product design, as shown in Figure 1), which effectively determined how the available bandwidth in each location was going to be split up into lots. Each lot is a particular frequency band covering a particular geographic area. In the case of the AWS auction, the FCC decided that six frequency blocks of paired spectrum (A-F) were to be auctioned, as shown above. Three blocks were 20 MHz and three were 10 MHz. Because the US is so large, each frequency block was also partitioned geographically. And because the FCC was attempting to accommodate all types of bidders, the FCC partitioned the blocks in three different ways: for blocks D-F the country was split into 12 large regions; for blocks B and C the country was split into 176 medium-size regions; and for block A the country was split into 734 small regions. Remarkably, the different partitions do not form a hierarchy in the sense that you cannot construct one of the medium-sized lots by aggregating a number of small lots. This inability to aggregate small into medium clearly limits substitution across blocks.

Figure 2. The absence of arbitrage across substitutes in the US AWS auction

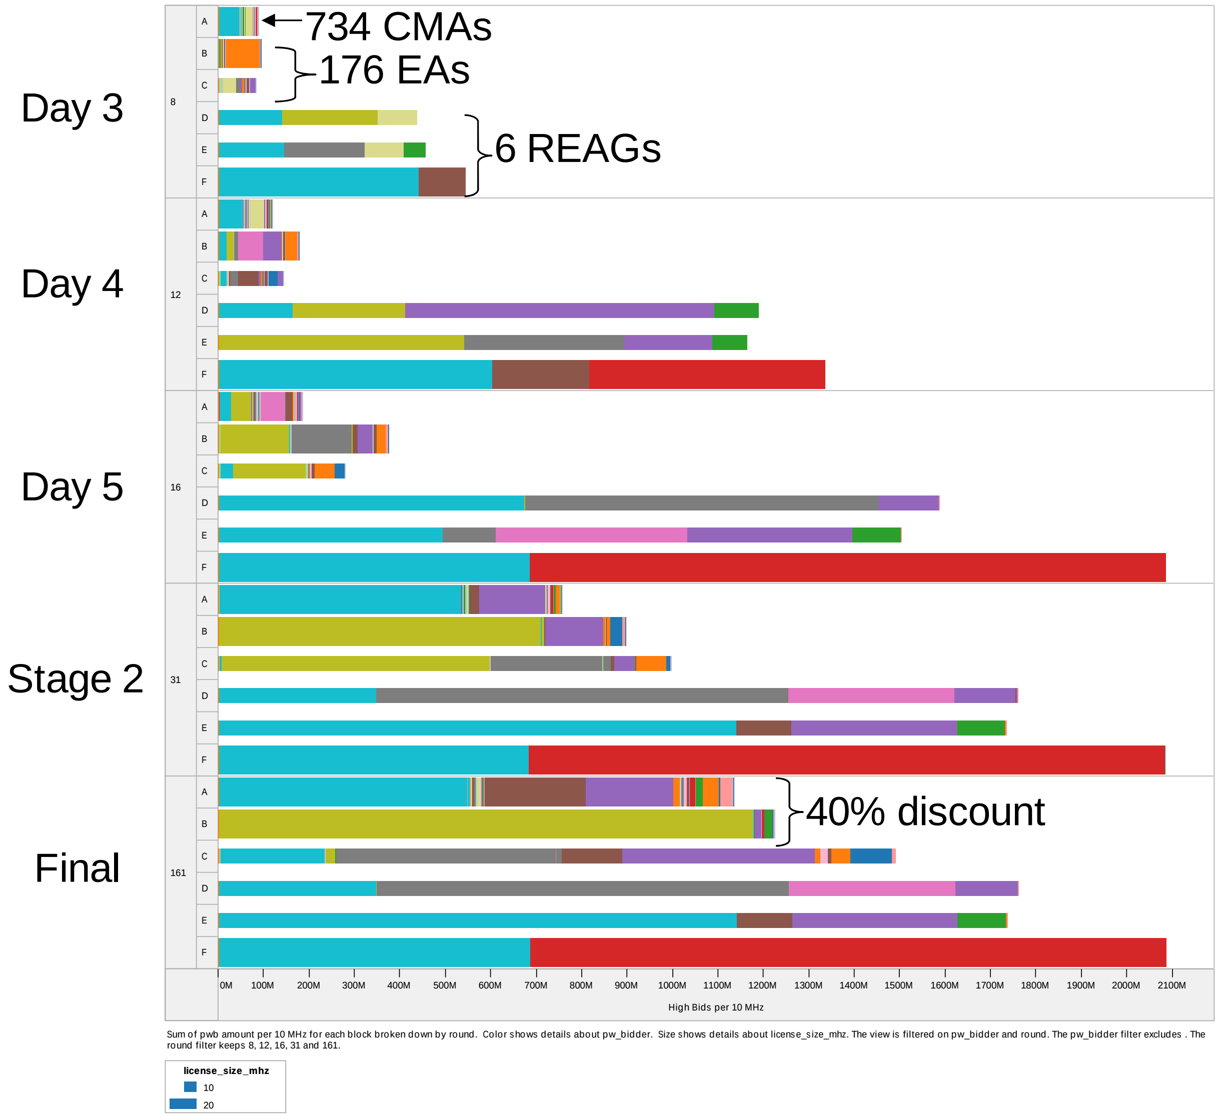

The underlying substitution problem was caused both by the product design—the use of specific blocks following three different geographic schemes—and the auction format. Figure 2 illustrates the severe problems bidders had substituting across blocks in the AWS auction. It shows the price per 10 MHz of spectrum for each of the blocks at the end of critical days in the auction. Recall there are six blocks, so there are six bars (A through F) at the end of each day. The 20 MHz bars are twice as thick as the 10 MHz bars so the area of the bar corresponds to revenues at the time indicated. Finally, different colors represent different bidders, so you can see who the provisional winners are at the various times in the auction. The two largest bidders are T-Mobile (turquoise) and Verizon (red).

If there was perfect arbitrage across blocks, then what you would see is the length of the bars would be the same at each time in the auction, indicating equal prices across blocks. Over time, the prices would move higher, but the prices would tend to move together across the blocks, as bidders would arbitrage to the cheaper lots per MHz of spectrum.

What happened in the AWS auction is extremely far from that. Look at the end of day five. At this point, the F block has already reached its final price. The A block is less than one twentieth the price of the F block. If the A block is roughly equivalent to the F block, why wouldn't Verizon, say, switch to the much cheaper A block, instead of placing bids twenty times higher on the F block? The reason has to do with substitution difficulties. When Verizon is bumped off a large F block license, it is easy for Verizon to substitute down to the A block, submitting say the 100 or so bids on the A lots that roughly cover the corresponding F lot. The problem is that once shifting down it would be nearly impossible to shift back up to F. The reason is that in subsequent rounds Verizon would only be bumped from some of the corresponding A block lots. Verizon would have to withdraw from many A lots in order to return to F, exposing itself to large withdrawal penalties. In addition on block A, Verizon would be vulnerable to various hold-up strategies, where speculators could pick important holes in a synergistic aggregation of lots.

Since substituting down from large (F, E, D) to small (C, B, A) lots is easier than substituting up, the auction essentially proceeded in a sequential fashion. First, the bidders competed for the large-lot blocks (F, E, D), then they competed for the medium-lot blocks (C and B), and finally the competition fell to the small-lot block (A). This explains the sequential, rather than simultaneous price process across blocks. See Bulow et al. (2009) for more on this auction.

Table 1. Band plan and final prices ($/MHz-pop) for paired spectrum in 700 MHz auction

Block	A	B	C
Bandwidth	12 MHz	12 MHz	22 MHz
Type	paired	paired	paired
Partition	176	734	12
Price	$1.16	$2.68	$0.76

The next major auction in the US was the 700 MHz auction. The band plan for the paired spectrum is shown above. The FCC did the same thing in this auction. Specific block s were auctioned, using three different partitions of the US. Again the different partitions did not form a hierarchy. The final prices per MHz-pop range from $0.76 for the C block to $2.68 for the B block, as shown in Table 1. These final prices differ by over a factor of three. We see again that the substitution across blocks is far from perfect. Interestingly, this time it is the small-lot block B that sold for a high price, and the large-lot block C that sold for a low price—just the opposite of what happened in the AWS auction.

Although the C-block had an open access provision, requiring that the operator not discriminate against either devices or applications, the terms of open access were sufficiently watered down that I doubt it had much of an impact on the C-block price. In my view, the price difference was because competing bidders thought that competing on the C-block against Verizon (or perhaps AT&T and Verizon) was sufficiently hopeless that it would be better to focus on the A and B blocks. See Cramton et al. (2007) for more on the competitive issues in this auction.

The conclusion from the now long history of spectrum auctions using the simultaneous ascending auction is that it works reasonably well in simple situations with a single geographic scheme. However in more complex settings, the approach leads to complex bidding strategies that complicate the auction and may undermine the efficient assignment of spectrum.

3 A better way: the package clock auction

Fortunately, there is a better way. All that is needed is a number of complementary enhancements that ultimately simplify the bidding process, improve its efficiency, and greatly expand its power.

First, much of the game playing, such as tacit collusion and other bid signaling, can be eliminated with a shift to anonymous bids. In a package clock auction the round-by-round revelation of information is limited to aggregate measures of competition. Limiting round reports to prices and excess demand for each product gives the bidders the information needed to form expectations about likely prices and in resolving common value uncertainty, yet such reports do not allow the signaling strategies that support tacit collusion. Moreover, the streamlined report simplifies bidder decision-making and keeps the bidders focused on what is most relevant, the relationship between prices and aggregate demand.

In most instances, spectrum lots covering the same region in adjacent frequencies are nearly perfect substitutes. To a close approximation, the bidder simply cares about the quantity of spectrum in MHz it has in the region, rather than the exact frequency location. Moreover, to minimize interference problems bidders prefer contiguous spectrum within any region. Thus, it makes sense in the initial stage to auction generic spectrum. The stage determines the quantity of contiguous spectrum won in each region. In this stage the spectrum is treated as if it were a homogenous good within each region. This is an enormous simplification of what is being sold. The idea is to treat each MHz of spectrum within a geographic region and a particular frequency band as perfect substitutes. We let the auction resolve first the main question of how much spectrum in each region each winner gets and at what price, before the auction turns to the more subtle and less important question of the exact frequencies.

Of course, there are some auctions where the differences across frequencies are too great to allow this simplified treatment, for example because of major interference differences by frequency, as the result of incumbents with a right to stay in the particular band. In such cases, the specific spectrum lots can be auctioned from the start, but in most cases, it is desirable to auction generic spectrum first and then determine the specific assignment in a second stage.

The specific assignment stage is simplified, since it only involves winners in the generic stage. The number of specific assignments typically is limited to the number of ways that the winners can be ordered. Thus, if there are m winners there are m! different specific assignments. For example, an auction with four winners in a particular region would have 4! = 4×3×2 = 24 different possible specific assignments. Assuming separability across regions, each of the four bidders would only need to express preferences among at most 24 different specific assignments. This number is reduced further if we assume that the bidder only cares about its own specific assignment and not the location of the other winners, as is commonly the case. Then for example with four winners of equal size, each winner would only need to express three preferences: the incremental value from the bidder's third-best specific assignment compared with its fourth-best, the incremental value from the bidder's second-best assignment compared with its third-best, and the incremental value from the bidder's first-best assignment compared with its second-best.

The use of generic lots, wherever possible, simplifies the auction, enhances substitution, and improves price discovery. Despite these advantages the FCC has chosen in each of its 72 auctions to sell specific lots. This is a common mistake in auction design. Interestingly, even in countries that recognized the advantages of selling generic lots, such as the German 3G auction, the generic lots were auctioned using a method for specific lots; that is, in the German 3G auction, even though the lots were perfect substitutes, the bidders bid on specific lots.

Once generic lots are adopted the next innovation becomes easier to see—the adoption of simple and powerful techniques that are well-suited to auctioning many divisible goods.

The first innovation is an improved product design, based on generic spectrum in each service area, accommodating multiple types of use.

The second innovation is the use of a simultaneous clock auction. This is a simplification of the simultaneous ascending auction. Each product has its own clock indicating its current price. Because of generic lots, each product may consist of multiple lots. In each round, the bidder is asked to indicate for each product the quantity of lots desired at the current price. At the end of the round, the auctioneer adds up the individual bids and reports the excess demand for each product. The price is then increased on any product with excess demand. This process is repeated until there is no excess demand on any product.

The two critical differences between the clock auction and the simultaneous ascending auction are 1) the bidder only answers demand queries, stating the quantities desired at the announced prices, and 2) there is no need to identify provisionally winning bidders at the end of every round.

The third innovation is more subtle, but extremely powerful. One can interpret the demand vector reported by each bidder in each round as a package bid. The bidder is saying, “At these prices, I want this package of lots.” Taking this interpretation seriously yields a combinatorial auction (or package auction) without the need for any optimization. This allows bidders to express complementarities within a simple price discovery process.

Larry Ausubel and I have been conducting exactly this sort of package auction since 2001 for electricity and gas products in France, Germany, Belgium, Denmark, Spain, Hungary, and the United States (Ausubel and Cramton 2004). Thus, far we have conducted over 60 high-stakes auctions with this format for assets worth about $10 billion. We also used the approach in a spectrum auction in Trinidad and Tobago in 2005. The approach has been highly successful.

The clock auction may end with some products in excess supply, as a result of complementarities among lots. In addition, since the clock process follows a single price path and only includes a limited number of price points, it is desirable to allow the bidder to specify additional bids in a supplementary round following the clock stage. The purpose is to let the bidder express preferences for additional packages that were missed by the clock process. In addition, the bidder can improve its bids on packages already bid on in the clock stage.

Once the clock bids and the supplementary bids are collected, an optimization is run to determine the value-maximizing generic assignment and prices. This two-step process of a clock auction followed by supplementary bids, which I call a package clock auction, was proposed by Larry Ausubel, Paul Milgrom, and me for spectrum auctions at an FCC auction conference in 2003 (Ausubel et al. 2006). We proposed the same approach for spectrum auctions in the UK in 2006, as well as for airport takeoff and landing rights in 2003. Meanwhile, Porter et al. (2003) demonstrate in the experimental lab the high efficiency of a closely related approach.

Two critical elements of a successful package clock auction are the pricing rule and the activity rule. I will discuss both at length. These two important rules work together to ensure that the bids are an accurate expression of bidder preferences throughout the entire auction. The high efficiency of the package clock auction derives mainly from incentives for nearly truthful bidding. A pricing rule based on second pricing encourages truthful bidding; whereas, the activity rule based on revealed preference ensures that these incentives for truthful bidding are felt throughout the clock stage.

4 UK spectrum auctions

The need for a technology neutral auction is commonplace in today’s world of rapidly developing communications technologies and applications. While the regulator can typically identify the viable candidate technologies based on early development, the regulator cannot decide how available spectrum should be split among the technologies without a market test. Examples are numerous, and several will be discussed here.

Ofcom, the independent regulator and competition authority for the UK communications industries, was the first to recognize and act on this need for a technology neutral auction. In spring 2006, Larry Ausubel and I proposed to Ofcom a version of the package clock auction. Since June 2006, I have been working with Ofcom in developing, testing, and implementing the design for a number of its auctions. Two such auctions, the 10-40 GHz auction and the L-band auction have occurred already. Both went very well, and provided a useful field test for the economically much larger 2.6 GHz auction scheduled to take place in the second-half of 2009. The Netherlands has also adopted this approach for its 2.6 GHz auction. Finally, Ofcom has proposed to use the package clock auction for the digital dividend auction to take place in 2010. The initial proposal presents some computational challenges given the quantity of spectrum and the number of technologies that are competing for that spectrum, although I suspect that some simplifications will be made as the range of possibilities narrows as we get closer to the auction date.

Ofcom has three main goals for the auction design. The auction should be technology neutral, allowing alternative viable technologies to compete for the spectrum on an equal basis. The auction should accommodate flexible spectrum usage rights, permitting the user to decide how the spectrum would be used, subject to minimizing interference externalities with neighbors. And the auction should promote an efficient assignment of the spectrum, putting the spectrum to its best use.

Simplicity and transparency are important secondary objectives. On simplicity, Ofcom recognized that satisfying the main objectives posed serious challenges, which could not be addressed with an auction design that is too simple. Moreover, simplicity has to be assessed recognizing the complexity of bidder participation. For example, the simultaneous ascending auction has simple rules, but incredibly complicated bidding strategies. In contrast, the package clock auction has more complex rules, but the rules have been carefully constructed to make participation especially easy. For the most part, the bidder can focus simply on determining its true preferences for packages it can realistically expect to win. In a package clock auction it is the auctioneer that needs to do the complex optimization, whereas the bidders can focus on their values for realistic packages.

Revenue maximization was explicitly excluded as an objective. Nonetheless, an efficient auction necessarily will generate substantial revenues. Indeed, my advice to countries is to focus on efficiency. A focus on revenues is short-sighted. In my view, the government is better off finding as much spectrum as possible and then auctioning it so as to put the spectrum to its best use. This approach creates a competitive and innovative market for communications, which has substantial positive spillovers to the rest of the economy. Under this approach, long-term revenues likely will far exceed those that would come from the maximization of short-term auction revenues.

4.1 The UK 2.6 GHz auction

The UK 2.6 GHz auction illustrates well the benefit of the package clock auction. This auction is for 190 MHz of spectrum at frequencies (2.6 GHz) that make it especially well-suited for mobile broadband communications. Currently, there are two fourth-generation (4G) technologies competing worldwide to be the next standard for broadband wireless communications: LTE, which uses paired spectrum, and WiMAX, which uses unpaired spectrum. Both technologies work well with a lot size of 5 MHz, but LTE requires that the lots come is pairs, one for uplink and another for downlink. In addition, a 5 MHz guard band is required between the paired and unpaired use, as well as between any two competing unpaired operators.

Figure 3. CEPT band plan from Electronic Communications Committee Decision (05)05

Type	Paired (FDD uplink)													
Lot	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Frequency	2500	2505	2510	2515	2520	2525	2530	2535	2540	2545	2550	2555	2560	2565

Type	Unpaired (TDD)									
Lot	1	2	3	4	5	6	7	8	9	10
Frequency	2570	2575	2580	2585	2590	2595	2600	2605	2610	2615

Type	Paired (FDD downlink)													
Lot	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Frequency	2620	2625	2630	2635	2640	2645	2650	2655	2660	2665	2670	2675	2680	2685

Confronted with this setting, the European Commission decided in 2005 that the band plan should be as shown in Figure 3 with 140 MHz of spectrum set aside for paired use and the remaining 50 MHz for unpaired use. Each 2×5 MHz paired lot would consist of an uplink on the lower block and a downlink on the upper block with 120 MHz separation between the two. This is an example of a specific band plan. If Ofcom adopted this plan, they would foreclose other possibilities that may put the spectrum to better use. In particular, the Commission had no way to know that splitting the spectrum 140 MHz paired and 50 MHz unpaired is best. Perhaps an alternative band plan with 100 MHz of paired spectrum as shown in Figure 4 would be better.

Figure 4. Increase in unpaired spectrum maintaining 120 MHz duplex spacing

Type	Paired (FDD uplink)										Unpaired (TDD)														Paired (FDD downlink)										Unpaired			
Lot	1	2	3	4	5	6	7	8	9	10	1	2	3	4	5	6	7	8	9	10	11	12	13	14	1	2	3	4	5	6	7	8	9	10	15	16	17	18
Frequency	2500	2505	2510	2515	2520	2525	2530	2535	2540	2545	2550	2555	2560	2565	2570	2575	2580	2585	2590	2595	2600	2605	2610	2615	2620	2625	2630	2635	2640	2645	2650	2655	2660	2665	2670	2675	2680	2685

Or perhaps the future of WiMAX is so bright that all of the spectrum should be unpaired as shown in Figure 5.

Figure 5. All unpaired spectrum

Type	Unpaired (TDD)																																					
Lot	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38
Frequency	2500	2505	2510	2515	2520	2525	2530	2535	2540	2545	2550	2555	2560	2565	2570	2575	2580	2585	2590	2595	2600	2605	2610	2615	2620	2625	2630	2635	2640	2645	2650	2655	2660	2665	2670	2675	2680	2685

Fortunately, a package clock auction can resolve the question of the best band plan at the same time that spectrum winners are identified and the spectrum is priced. There is no reason that the regulator needs to specify the band plan; the auction can do that based on the competitive bids of the operators. Indeed, this can be an accomplished with an auction that is much simpler for bidders than the simultaneous ascending auction.

Here is how it works. The 140 MHz is split into 38 lots, 5 MHz each. The lots are perfect substitutes, since they are generic. All bids are package bids, so there is no exposure problem.

In the clock stage, the bidder simply specifies how many paired lots and how many unpaired lots it desires at the current prices. Packages of paired lots are guaranteed to be contiguous, as are packages of unpaired lots. In order to achieve 120 MHz spacing, consistent with the EC decision, it is necessary to have a minimum of 9 unpaired lots plus one 5 MHz guard block between the paired and unpaired use. If demand for paired and unpaired spectrum is such that this minimum constraint is not binding, then the paired and unpaired prices are linked at 2-1: one 2x5 MHz paired lot is priced at two times an unpaired lot; otherwise, if as a result of the minimum constraint of 9 unpaired lots, the price of unpaired lots may be lower than the paired price per 5 MHz. Thus, the clock stage has at most two prices, the price for unpaired lots and the price for paired lots. So long as there is excess demand of lots the clock prices continue to increase. Indeed, the clock prices will continue to increase until the point that the demands are feasible, which means that the demands can be met with contiguous awards. Notice that the clock stage is much simpler than the simultaneous ascending auction.

Following the clock stage, the bidders submit supplementary bids. The supplementary bids are either improvements to clock bids or bids on additional packages that were not bid on in the clock stage. For unpaired bidders, these additional bids can include the possibility of an award split between the lower block of unpaired spectrum and an upper block, which would be necessary if more than 9 unpaired lots are awarded.

Once the clock and supplementary bids are collected, the auction system takes all these bids and performs a series of optimizations to determine the value maximizing generic assignment, and the base prices to be paid by each winner.

The system also determines the feasible options each winner faces for specific assignments. These are reported to each winner, and the winner is asked to submit assignment bids indicating the incremental value derived from each option that is better than the winner's least preferred option. These assignment stage bids are then collected and another series of optimizations is performed to determine the value maximizing specific assignments and the additional payments the winners make in addition to the base prices calculated at the end of the supplementary round.

I now explain the details of two essential rules in the package clock auction: the pricing rule and the activity rule. The rules may appear complex, but the complexity actually simplifies the bidding strategies, making it easier for bidders to participate in the auction.

5 Pricing rule: Closest-to-Vickrey core pricing

Prices are determined at two points in the auction, after the clock stage, including the supplementary bids, to determine the base prices for the winners in the value-maximizing generic assignment, and after the assignment stage to determine the additional payments for specific assignments.

The pricing rule plays a major role in fostering incentives for truthful bidding. Pay-as-bid pricing in a clock auction or a simultaneous ascending auction creates incentives for demand reduction (Ausubel and Cramton 2002). Large bidders shade their bids, recognizing their impact on price. This bid shading both complicates bidding strategies and also leads to inefficiency.

In contrast, Vickrey pricing provides ideal incentives for truthful bidding. Each winner pays the social opportunity cost of its winnings, and therefore receives 100 percent of the incremental value created by its bids. This aligns the maximization of social value with the maximization of individual value for every bidder. Thus, with private values, it is a dominant strategy to bid truthfully. See Ausubel (2004, 2006) for an analysis in a clock auction.

Unfortunately, as a result of complements, it may be that the Vickrey prices are too low in the sense that one or more bidders would be upset with the assignment and prices paid, claiming that they had offered the seller more. For example, suppose there are two items, A and B. Bidder 1 bids $4 for A, bidder 2 bids $4 for B, and bidder 3 bids $4 for A and B. The Vickrey outcome is for 1 to win A, 2 to win B, and each winner pays $0. Bidder 3 in this case has a legitimate complaint, "Why are you giving the goods to bidder 1 and 2, when I am offering $4 for the pair?" The basic problem is that with complements, the Vickrey outcome may not be in the core. Some coalition of bidders may have offered the seller more than the sum of the Vickrey prices. This point has been emphasized in Ausubel and Milgrom (2002). (The core is defined as a set of payments that support the efficient assignment in the sense that there does not exist an alternative collation of bidders that has collectively offered the seller more.)

The solution is to increase one or more prices to assure that the prices are in the core. In order to provide the best incentives consistent with core pricing, we find the lowest payments that are in the core; that is, such that no alternative coalition of bidders has offered the seller more than the winning coalition is paying.

If we are auctioning a single item, then this is the second-price auction. Suppose the highest bidder bids $100 and the second-highest bidder bids $90. The item is awarded to the highest bidder, who pays the second-highest price of $90—the social opportunity cost of awarding the good to the highest bidder. Alternatively, we can think of assigning the item to maximize value, so we assign it to the highest bidder, and then we find the smallest payment that satisfies the core constraints. In this case, the second-highest bidder would be upset if the highest bidder paid less than $90, so $90 is the bidder-optimal core price. When the items are substitutes, then the bidder-optimal core point is unique and identical to the Vickrey prices.

Typically, the payment minimizing core prices, or bidder-optimal core prices, will not be unique. Thus, it will be important to have a method of selecting a unique bidder-optimal core point when there are many. The sensible approach adopted in each of the recent Ofcom auctions for both the base prices and the assignment prices is to select the payment minimizing core prices that are closest to the Vickrey prices. This is what I call closest-to-Vickrey core pricing. Since the set of core prices is convex—a polytope formed from the intersection of half-spaces—and the Vickrey prices are always unique, there is a unique vector of core prices that is closest in Euclidean distance to the Vickrey prices. Not only are the prices unique, but since they are bidder-optimal-core prices, they maximize the incentive for truthful bidding among all prices that satisfy core constraints.

The approach then is to take all the bids from the clock stage and the supplementary bids, determine the value maximizing assignment, and then determine the payment minimizing core prices that are closest to the Vickrey prices. It is my experience that bidders are quite happy with this approach—they like the idea of minimizing payments, and they recognize the importance of making sure that the prices are sufficiently high that no coalition of bidders has offered the seller more. Prices are as small as possible subject to all the competitive constraints.

Calculating the winning assignments and prices involves solving a sequence of standard optimization problems. The basic problem is the winner determination problem, which is a well understood set-packing problem. The main winner determination problem is to find the value maximizing assignment. To guarantee uniqueness, there is a sequence of lexicographic objectives, such as: 1) maximize total value, 2) minimize concentration, 3) maximize quantity sold, and 4) random. Thus, first you maximize total value. Then you add a constraint that the value equals this maximum value and you minimize concentration. Then you add another constraint that the concentration equals this minimum concentration and you maximize the quantity sold. Finally, you add a constraint that the quantity sold equals this maximum quantity and you maximize an objective based on random values for each bid rather than the true bids and resolve the optimization. This guarantees uniqueness.

Calculating the prices is a bit more involved. First, we determine the Vickrey prices by solving a sequence of winner determination problems, essentially removing one winner at a time to determine each winner's social opportunity cost of winning its package. Then we determine the bidder-optimal core prices using a clever constraint generation method proposed in Day and Raghavan (2007). Having found the Vickrey prices, we solve another optimization to find the most violated core constraint. If there is none, then we are done, since the Vickrey prices are in the core. Otherwise, we add this most-violated constraint and resolve the optimization, again finding the most violated core constraint. We add it to the optimization and re-solve. We keep doing this until there is no violated core constraint, and then we are done.

The reason that that Day-Raghavan approach is a highly efficient method of solution is because in practice there are typically only a handful of violated core constraints; thus, the procedure stops after just a few steps. In contrast the number of core constraints grows exponentially with the number of bidders and that makes including all the core constraints explicitly an inefficient method of solving the problem, both in time and memory.

As mentioned, the tie-breaking rule for prices is going to be important, since typically ties will arise along the southwest face of the core polytope. Finding the prices that are closest to the Vickrey prices involves solving a simple quadratic optimization. This gives us a unique set of prices. Uniqueness is important; it means that there is no discretion in identifying the outcome, either in the assignment or the prices.

An example will help illustrate all of these concepts. Suppose there are five bidders, 1, 2, 3, 4, 5, bidding for two lots, A and B. The following bids are submitted:

$b_1\{A\} = 28$

$b_2\{B\} = 20$

$b_3\{AB\} = 32$

$b_4\{A\} = 14$

$b_5\{B\} = 12$

Bidders 1 and 4 are interested in A, bidders 2 and 5 are interested in B, and bidder 3 is interested in the package A and B.

Determining the value maximizing assignment is easy in this example. Bidder 1 gets A and bidder 2 gets B, generating 48 in total value. No other assignment yields as much. Vickrey prices are also easy to calculate. If we remove bidder 1, then the best assignment gives A to bidder 4 and B to bidder 2, resulting in 34, which is better than the alternative of awarding both A and B to bidder 3, which yields 32. Thus, the social opportunity cost of bidder 1 winning A is 34 – 20 = 14 (the value lost from bidder 4 in this case). Similarly, if we remove bidder 2, then the efficient assignment is for bidder 1 to get A and bidder 5 to get B, resulting in 40. Then the social opportunity cost of bidder 2 winning B is 40 – 28 = 12 (the value lost from bidder 5). Hence, the Vickrey outcome is for bidder 1 to pay 14 for A and for bidder 2 to pay 12 for B. Total revenues are 14 + 12 = 26. Notice that bidder 3 has cause for complaint, since bidder 3 offered 32 for both A and B.

Now consider the core for this example. The core is represented in the payment space of the winning bidders—in this case the payments of bidders 1 and 2. Each bid defines a half-space of the payment space:

- Bidder 1's bid of 28 for A implies 1 cannot pay more than 28 for A.
- Bidder 2's bid of 20 for B implies 2 cannot pay more than 20 for B.
- Bidder 3's bid of 32 for AB implies that the sum of the payments for A and B must be at least 32.
- Bidder 4's bid of 14 for A implies that bidder 1 must pay at least 14 for A.
- Bidder5's bid of 12 for B implies that bidder 2 must pay at least 12 for B.

The core is the intersection of these half-spaces as shown in Figure 6.

Figure 6. The Core

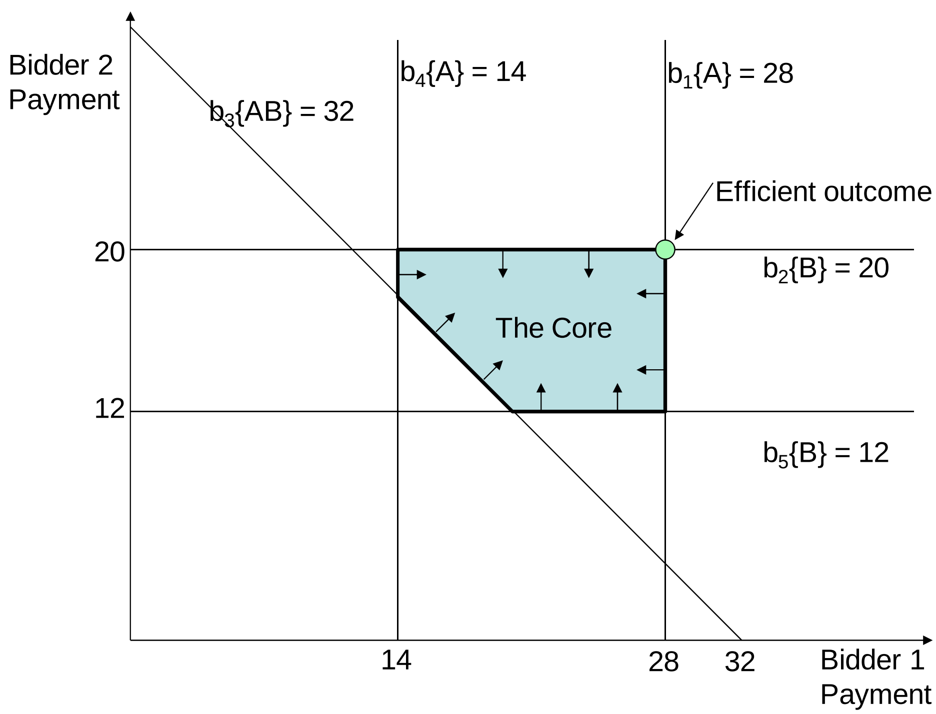

This example is quite general. First, unlike in some economic settings, in an auction, the core is always nonempty. The reason is that the core always includes the efficient outcome. The reason is that all the constraints are southwest of the efficient point, since the efficient point maximizes total value. Second, the core is always a convex polytope, since it is the intersection of numerous half-spaces. Third, complementarities, like bidder 3's bid for AB, are the source of the constraints that are neither vertical nor horizontal. These are the constraints that can put the Vickrey prices outside the core. Without complementarities, all the constraints will be vertical and horizontal lines, and there will be a unique extreme point to the southwest: the Vickrey prices.

Figure 7. Vickrey prices: how much can each winner's bid be reduced holding others fixed?

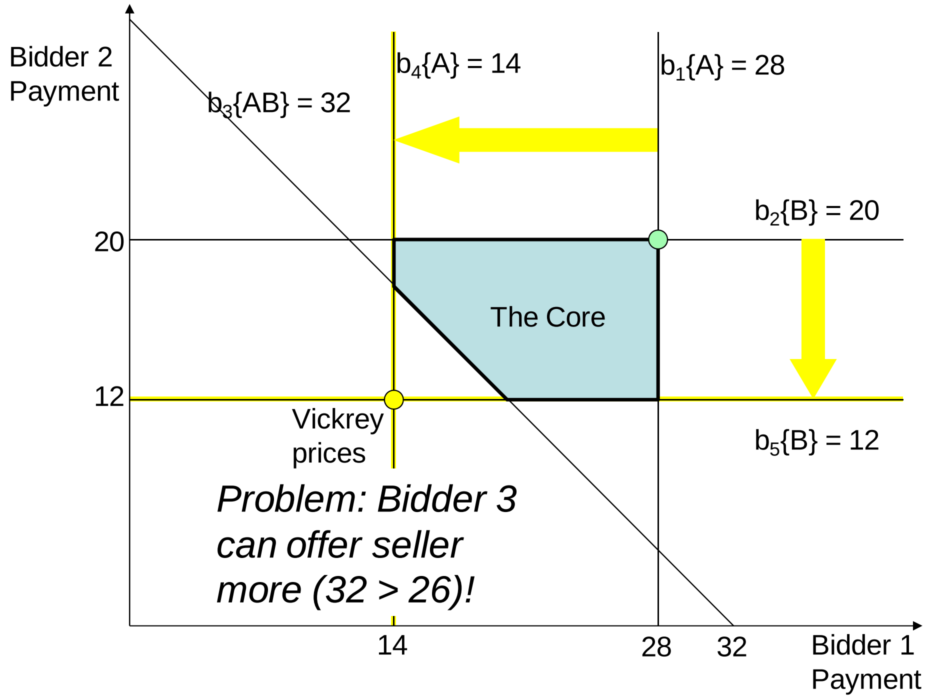

The graphical representation of the core is also a useful way to see the Vickrey prices. Vickrey is asking how much can each winner unilaterally reduce its bids and still remain a winner. As shown in Figure 7, bidder 1 can reduce its bid to 14 before bidder 1 is displaced by bidder 4 as a winner. Similarly, bidder 2 can reduce its bid to 12 before being displaced by bidder 5. Thus, the Vickrey prices are 14 and 12. The problem is that these payments sum to 26, which violates the core constraint coming from bidder 3's bid of 32 for AB.

Figure 8. Bidder-optimal core prices: jointly reduce winning bids as much as possible

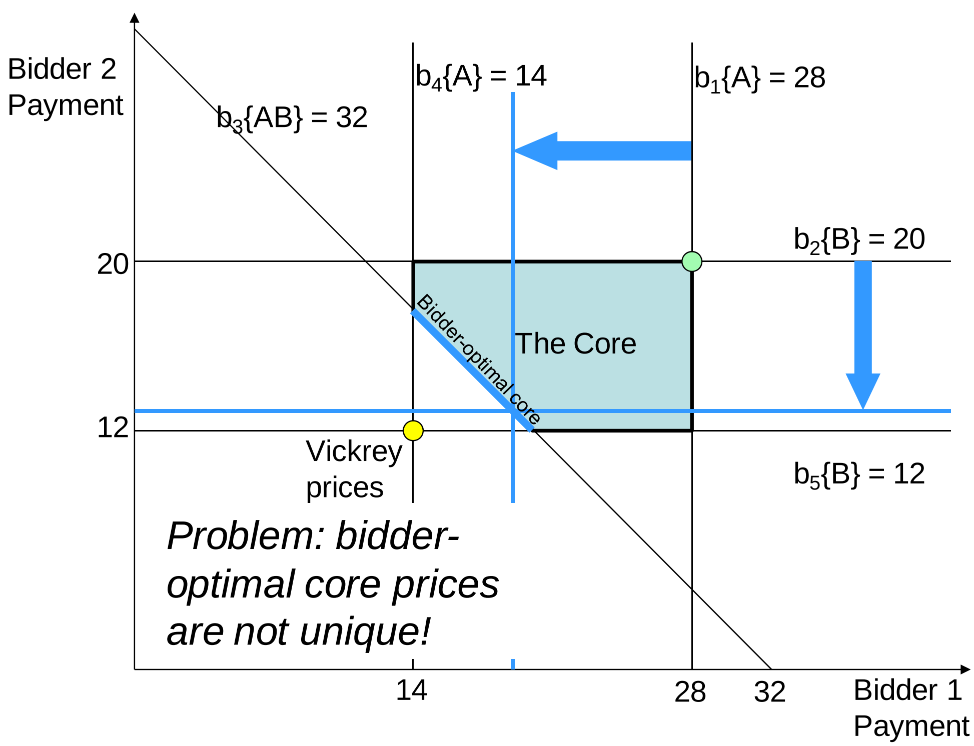

Bidder-optimal core prices can also be thought of maximal reductions in the bids of winners, but rather than reducing the bids of each winner one at a time, we jointly reduce all the winning bids, as shown in Figure 8, until the southwest face of the core is reached. As can be seen, this does not result in a unique core point, since the particular point on the southwest face depends on the rate at which each winner's bids are reduced. The bidder-optimal core points consist of the entire southwest face of the core. If the southwest face is a unique point, then it is the Vickrey prices; if the southwest face is not unique then the face is a core constraint involving complementarities, and the Vickrey prices lie outside the core.

Figure 9. Core point closest to Vickrey prices

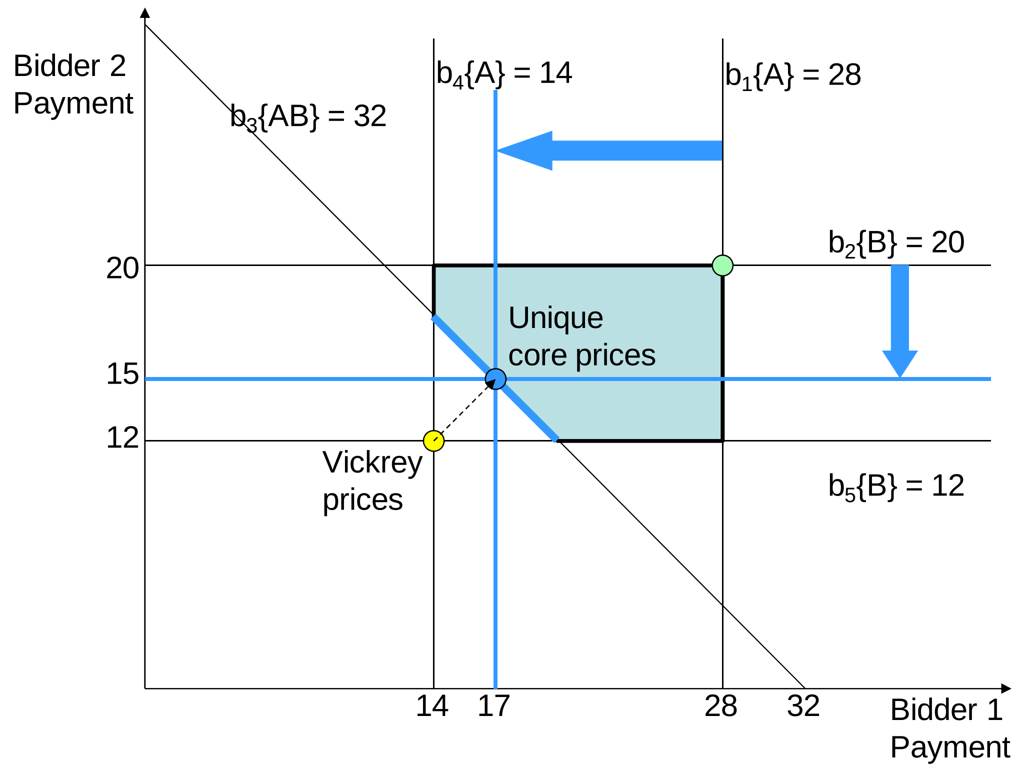

Nonetheless, there is always a unique core point that is closest to the Vickrey prices. This is seen in Figure 9, as the bidder-optimal core point that forms a 90 degree angle with the line that passes through the Vickrey prices. This point minimizes the Euclidean distance from the Vickrey prices.

Closest-to-Vickrey core pricing was adopted in each of the UK spectrum auctions, both the two that have already been held as well as the proposed auctions for the 2.6 GHz spectrum and the digital dividend spectrum. Closest-to-Vickrey core pricing is also proposed in the Netherlands 2.6 GHz auction, and in the US auction for takeoff and landing slots at the New York City airports.

Bidder-optimal core pricing has several advantages. First, it minimizes the bidders' incentive to distort bids in a Pareto sense: there is no other pricing rule that provides strictly better incentives for truthful bidding. Bidder-optimal core pricing implies Vickrey pricing, whenever Vickrey is in the core. For example, when lots are substitutes, Vickrey is in the core, and the bidders have an incentive to bid truthfully. Since the prices are in the core, it avoids the problem of Vickrey prices being too low as a result of complements. Finally, the rule has the desirable property that revenue is monotonic in bids and bidders. Adding either bids or bidders can only increase revenues. In contrast, Vickrey prices, as a result of complements, can either increase or decrease as either bids or bidders are added.

6 Activity rule: simplified revealed preference

Good price discovery is essential in realizing the benefits of a dynamic auction. Good price discovery stems from providing incentives for the bidders to make truthful bids throughout the auction process. The pricing rule discussed in the prior section is one essential element, but you also have to be concerned about what we see on eBay every day: bid sniping—jumping in at the last instance in an auction, holding your information back and not revealing it to the market. Bid sniping is an effective strategy in eBay auctions, and is typically used by experienced eBay bidders. The strategy is made possible on eBay, because eBay does not have an activity rule. In contrast, nearly all high-stake auctions, such as the FCC spectrum auctions, do have an activity rule. The FCC uses a quantity-based rule. This rule has worked reasonably well in the FCC's simultaneous ascending auctions, but in a package clock auction with closest-to-Vickrey core pricing, we need a more complex rule, one that is based on revealed preference. A rule based on revealed preference is effective at getting bidders to focus on their valuations and to bid in a profit maximizing way throughout the auction.

Absent an activity rule, bidders will have an incentive to hold back to conceal information. This bid sniping behavior is so common in eBay auctions that the auctions are better modeled as sealed-bid second-price auctions, rather than ascending auctions. The activity rule is intended to promote truthful bidding throughout the auction process. An effective activity rule will enhance price discovery, enabling the bidders to better focus on relevant packages and to resolve common-value uncertainty.

As mentioned, the traditional activity rule in both simultaneous ascending auctions and clock auctions has been a quantity-based rule: to be a large winner at the end of the auction, the bidder must be a large bidder throughout the auction. In particular, each lot corresponds to a particular quantity of spectrum, measured in either MHz-pop (the bandwidth times the population) or in eligibility points. The bidder starts with an initial eligibility based on the bidder's initial deposit. To maintain this level of eligibility in future rounds, the bidder needs to bid on a sufficiently large quantity of spectrum in the current round, where "sufficiently large" is stated as some percentage, typically between 80% and 100% of the bidder's current eligibility. If the bidder bids on a smaller quantity, the bidder's eligibility is reduced in future rounds. This quantity-based rule has worked reasonably well, although as mentioned, it does create an incentive for parking eligibility on lots that a bidder is not truly interested in, especially if the eligibility points are not a good measure of relative value across lots. (The FCC's MHz-pop measure is especially poor with small lots. Spectrum in New York City is much scarcer than spectrum in Montana. As a result, spectrum prices (and values) are much higher in New York City on a per MHz-pop basis. Despite this obvious fact, demonstrated in many dozens of spectrum auctions, the FCC still continues to use MHz-pop as the quantity measure in its auctions, exacerbating parking and other problems associated with the activity rule.)

In many clock auctions, an activity requirement of 100% is used, which means that the bidder cannot increase the size of the package, as measured in eligibility points, as prices rise. For the case of a single product, this means that the bidder must bid in a manner consistent with a downward-sloping demand curve.

In a package clock auction, one can use this quantity-based rule in the clock stage, but one also needs to specify how the rule limits bids in the supplementary round. This linkage between the clock bids and the supplementary bids is of critical importance, for otherwise the bidder could bid snipe, submitting all of its bids in the supplementary round.

Ofcom proposed the following, which I call the *eligibility point rule*. During the clock stage, the bidder cannot increase the package size. Moreover, whenever the bidder reduces the package size, the bid on all larger packages is capped by the prices at the time of the reduction. For example, if during the clock stage a bidder drops from a package of size 10 to 6 at prices p, then for all packages q of size 7 to 10, the supplementary bid cannot be more than $p \cdot q$.

The eligibility point rule, which Ofcom used in its first two package clock auctions, has the advantage of simplicity. For each package there is at most a single linear constraint on the supplementary bid. However, it has a potentially serious problem. The straightforward strategy of bidding on the most profitable package in the clock stage is a poor strategy. A bidder following such a strategy would find that its supplementary bids would be sharply constrained, well below true values. To avoid this problem, the bidder must instead bid in the clock stage to maximize package size, subject to a nonnegative profit constraint. That is, the bidder throughout the clock stage bids on the largest package that is still profitable.

Larry Ausubel, Paul Milgrom, and I proposed an alternative activity rule based on revealed preference for the package clock auction (Ausubel et al. 2006). Revealed preference is the underlying motivation for all activity rules. We wish to require that a bidder bid in a way throughout the auction that is consistent with the bidder's true preferences. Since we do not know the bidders true preferences, the best we can hope for is for the bidder to bid in a manner that is consistent with its revealed preferences. In the simplest case of a single-product clock auction, this is equivalent to monotonicity in quantity, just like the eligibility point rule, but when we have multiple products the two rules differ in important ways.

For the package clock auction, the *revealed preference rule* is as follows. (See Harsha et al. 2009 for a stronger statement.) During the clock stage, a bidder can only shift to packages that have become relatively cheaper; that is, at time $t' > t$, package $q_{t'}$ has become relatively cheaper than q_t:

(P) $$q_{t'} \cdot (p_{t'} - p_t) \leq q_t \cdot (p_{t'} - p_t).$$

Moreover, every supplementary bid $b(q)$ must be less profitable than the revised package bid $b(q_t)$ at t:

(S) $$b(q) \leq b(q_t) + (q - q_t) \cdot p_t.$$

That is, each clock bid for package q_t, as improved in the supplementary round, imposes a cap on the supplementary bid for package q.

An important advantage of the revealed preference rule is that a bidder following the straightforward strategy of bidding on its most profitable package in the clock stage would retain the flexibility to bid its full value on all packages in the supplementary round.

To illustrate the implications of the two activity rules, consider the following example with two bidders and two identical lots (one product) in a setting of substitutes. The bidders' preferences are given in the Table 2, indicating the marginal and average value for 1 lot and 2 lots.

Table 2. An example with two bidders and two identical lots

	Marginal Value		Average Value	
	Bidder A	Bidder B	Bidder A	Bidder B
1 lot	16	8	16	8
2 lots	2	2	9	5

Since the lots are substitutes, both bidders want to bid their true values in the supplementary round. However, consider what happens in the clock stage in response to the two different rules.

With the revealed preference rule, each bidder has an incentive to bid on its most profitable package in each round. Thus, the bidding simply moves up each bidder's marginal value (demand) curve. When the clock price reaches 2, both bidders drop from a package of size 2 to 1, and excess demand drops to zero. The clock stage ends at the competitive equilibrium price of 2 and the efficient assignment. Indeed, there is no need for any supplementary bids in this case. Bidder A can enter supplementary bids of 16 and 18, and bidder B can enter supplementary bids of 8 and 10, but these supplementary bids will not change the outcome in any way. Each bidder wins one lot and pays 2 (the Vickrey price). The supplementary round is unnecessary. The clock stage, by revealing the bidders marginal value information, up to the point of no excess demand, has revealed all that is needed to determine and price the efficient assignment. This is a general result with substitutes.

With the eligibility point rule, bidders are forced to distort their bidding away from the straightforward strategy of profit maximization. In order to preserve the ability to bid full values in the supplementary round, the bidders instead bid on the largest package that is still profitable. This entails moving up the average value curve, since when the average value is exceeded a package is no longer profitable. Thus, when the clock price reaches 5, bidder B's average value for 2 is reached and the bidder drops its demand to 1. Then when the clock price reaches 8, bidder B's average value for 1 is reached and bidder B drops out. At this point there is no excess demand, so the clock stage ends with bidder A demanding 2, bidder B demanding zero, and the clock price at 8. In the supplementary bid round, the bidders again submit their true preferences, and the optimization determines that each bidder should win one lot and should pay 2. The supplementary round was required to determine the efficient assignment and price the goods. Notice that the clock stage did little but mislead the bidders into thinking that bidder A would win all the items at a high price.

The reader might think that I somehow rigged this example to make the eligibility point rule look bad. This is not the case. Whenever lots are substitutes, the same features will be observed. With revealed preference, the clock stage will converge to the competitive equilibrium, revealing the efficient outcome and supporting prices; whereas with the eligibility point rule, the clock stage ends with an assignment that is excessively concentrated and prices that are too high. This result follows from the simple fact that average value exceeds marginal value, whenever aggregate demand is downward sloping, as shown in Figure 10. Having participated in many dozens of major spectrum auctions, I can confirm that this is indeed the typical case.

Figure 10. Downward sloping aggregate demand implies average value > marginal value

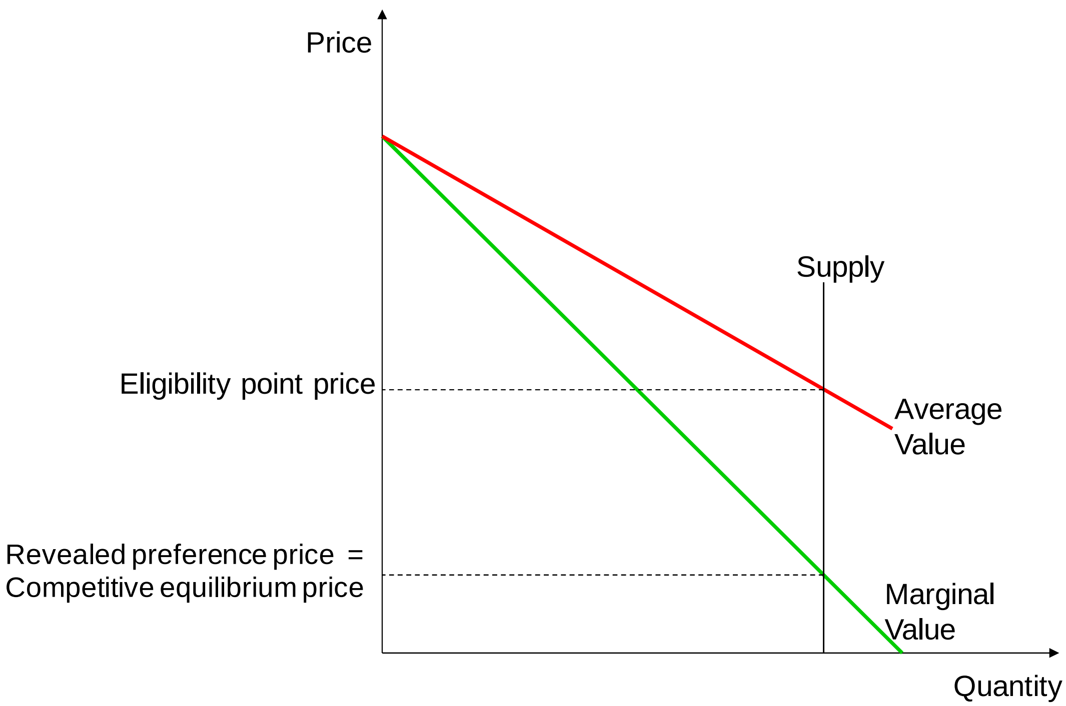

Figure 11 shows another example with a supply of 18 identical lots and five bidders, each with a different constant elasticity of demand. Marginal values are shown in the first five columns and average values are shown in the next five. The clock stage with the revealed preference rule is shown in green. The bidders bid to maximize profit and so reduce demands as marginal values are reached. The clock stage ends at a price of 370, the competitive equilibrium price and the assignment is efficient.

Figure 11. An example with five bidders and constant demand elasticity

Weaker bidder reveals too much

	Demand Elasticity									
	0.50	0.60	0.70	0.80	0.90					
	Marginal Value					Average Value				
Lots	Bidder A	Bidder B	Bidder C	Bidder D	Bidder E	Bidder A	Bidder B	Bidder C	Bidder D	Bidder E
1	10,000	4,642	2,683	1,778	1,292	10,000	4,642	2,683	1,778	1,292
2	2,500	1,462	997	748	598	6,250	3,052	1,840	1,263	945
3	1,111	744	558	450	381	4,537	2,282	1,413	992	757
4	625	461	370	314	277	3,559	1,827	1,152	823	637
5	400	317	269	238	216	2,927	1,525	975	706	553
6	278	234	207	189	176	2,486	1,310	847	620	490
7	204	181	166	156	149	2,160	1,149	750	553	441
8	156	145	138	132	128	1,909	1,023	674	501	402

Too concentrated

Bidding norm	Clearing Price	
Max profit	370	
Max size	1,292	Price too high
Difference	249%	

Supply	18
Misassigned	5
Fraction misassigned	28%

	Total Value
Max profit	31,428
Max size	29,150
Inefficiency	7.3%

In contrast, with the eligibility point rule, the outcome is shown in yellow. The clock stage outcome is too concentrated, misassigning 5 of the 18 lots and causing an inefficiency of 7.3%. The final price is 2.5 times the competitive equilibrium. The rule also forces the weakest bidder to reveal its entire demand schedule; whereas the strongest bidder reveals nothing.

What is essential for price discovery is the revelation of the marginal value information. This helps bidders make the marginal tradeoffs that are of greatest relevance in figuring out what the outcome should be. This is why I believe the eligibility point rule is a poor choice.

To further test the two activity rules, I conducted numerous simulations using realistic demand scenarios with significant complementarities from both technological and minimum scale constraints. I assumed that the bidders bid on the most profitable package with revealed preference (max profit) and bid on the largest profitable package with the eligibility point rule (max size). The results are summarized in Figure 12. Notice how the revealed preference rule achieves substantially higher efficiency in many fewer rounds.

Figure 12. Revealed preference rule yields higher efficiency and fewer bids in the clock stage

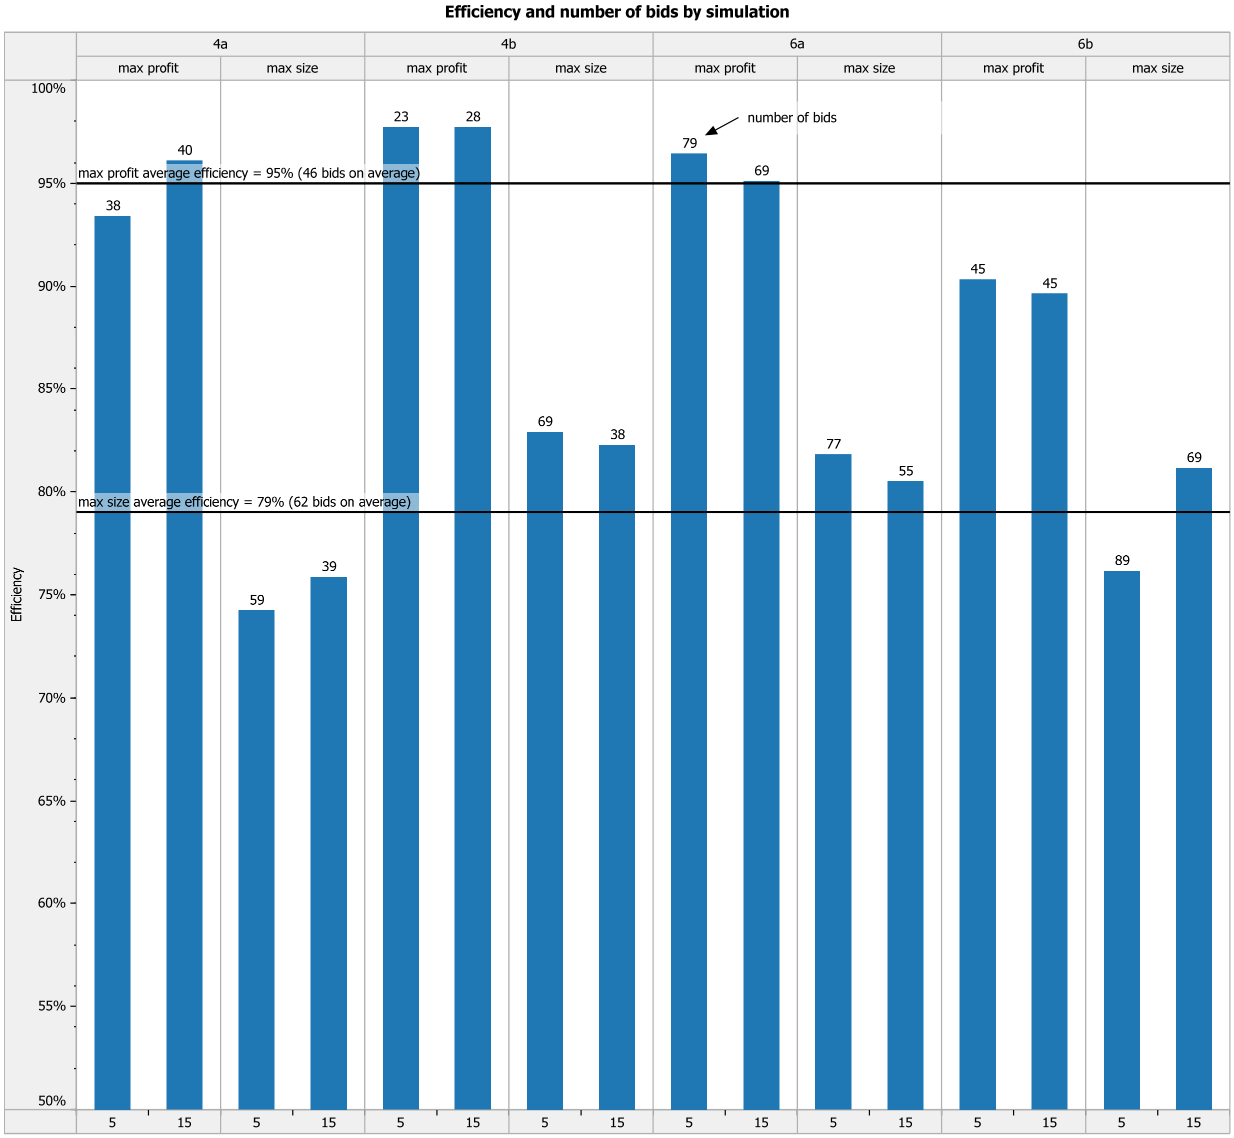

5 = low bid increments (5 to 15%); 15 = high bid increments (15 to 30%).

Figure 13 shows the better price discovery of the revealed preference rule as observed in the simulations. The top-two rows show the price path with a minimum bid increment of 5% under the revealed preference rule and the eligibility point rule, respectively. Notice the tendency for the price path under max-size bidding to go up too quickly and overshoot the competitive equilibrium price, creating excess supply (in red). The bottom two rows show the same scenarios, but with a larger minimum bid increment of 15%.

Figure 13. Price path for revealed preference and eligibility point rules by product

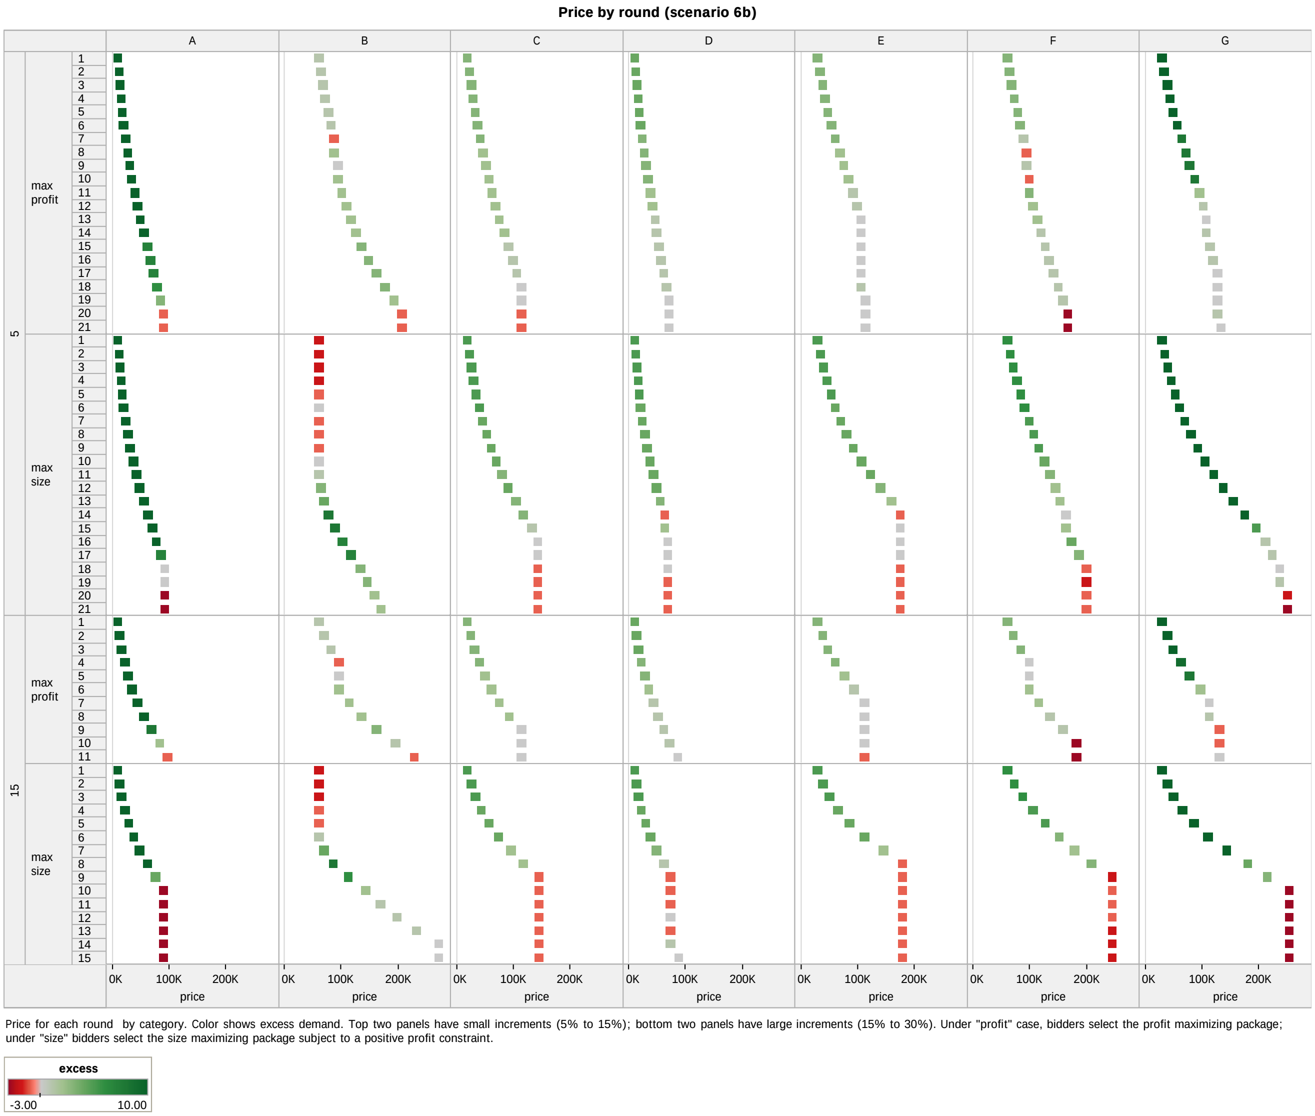

Price for each round by category. Color shows excess demand. Top two panels have small increments (5% to 15%); bottom two panels have large increments (15% to 30%). Under "profit" case, bidders select the profit maximizing package; under "size" bidders select the size maximizing package subject to a positive profit constraint.

As a final test of the two activity rules, as well as other elements of the auction design, I conducted a series of full-scale tests in the experimental lab. For the tests, the Ofcom auction platform was used and indeed Ofcom staff served as the auctioneer. The subjects in the test were PhD students, who had taken an advanced course in game theory and auction theory, and had prior participation in package clock auction experiments. I chose such an experienced and expert subject pool, since in the actual spectrum auctions bidders often hire experts and devote substantial time and money to understand the strategic implications of the rules.

Each subject participated in several auctions over a two-week period. In each auction, the subject was given a bidding tool, which calculated the subject's value for each package consistent with the bidder's business plan. The scenarios as represented by the various bidding tools were chosen to be realistic. The valuation models included both substitutes and complements. Complements came from minimum scale constraints as well as technological requirements. A training session was held before the auctions to explain the details of the package clock auction, including the two different activity rules. All subjects participated in both activity rule treatments. Each subject was paid an amount based on her experimental profits. The average subject payment was $420.

The experiments confirmed that the eligibility point rule caused a major deviation from straightforward bidding in the clock stage. Bidders quickly realized the need to bid on the largest profitable package. This undermined price discovery, but given the private value setting and simple valuation models, the poor performance of the clock stage was largely corrected by the supplementary bids and the optimization that followed. There were some instances of inefficiency when bidders deviated from bidding on the largest profitable package and then found they were unable to bid full values in the supplementary stage.

In contrast, with the revealed preference rule, bidders almost always followed the straightforward strategy of bidding on the most profitable package. In the supplementary round, bidders typically bid full value and were not constrained by the revealed preference rule. As a result, efficiency was nearly 100%.

One issue that was discovered in the lab was the complexity of the revealed preference rule. The few bidders who deviated from bidding on the most profitable package in each round of the clock stage found they were unable to bid full value in the supplementary round as a result of the revealed preference constraint. The difficulty for these bidders was that it was difficult for them to determine how high they could bid, since revealed preference is not a single constraint, but one constraint for each clock bid. Moreover, improving the bids on various clock packages causes the constraints to change. The challenge for the bidder is to figure out how best to adjust numerous bids in order to satisfy many constraints (one per round). Even the brightest PhD students found this to be a daunting task without some computational help.

One solution to the complexity problem is for the auction system to provide the bidder with some help. For example, the bidder could provide the system with its desired bids. The auction system then would indicate a summary of the bids that currently violate revealed preference constraints and suggest an alternative set of bids that satisfies all constraints and is closest (in Euclidean distance) to the desired bids. This is exactly the information the subjects in the lab were looking for in the few instances of deviations from straightforward bidding. In the lab, the deviations were minor and the bids would have been easily adjusted with the help of a smart auction system. The optimization I propose is a quadratic optimization with linear constraints. It is easily solved in an instant.

In addition to complexity, the revealed preference rule may at times be too strong. Bidders' values may change over the course of the auction for example as the result of common value uncertainty.

6.1 *Simplifying revealed preference*

A simplified revealed preference rule may address both the complexity and changing values issues. The idea behind the rule is that it may be unnecessary to include all of the revealed preference constraints to get the bidders to adopt the straightforward bidding. Since the incentive for bid sniping is not too strong, even the possibility of a revealed-preference constraint may be sufficient to induce the desired behavior. People put coins in parking meters in order to avoid the possibility of a parking ticket. We can hope that a simplified revealed preference rule will have the same effect in the package clock auction.

The *simplified revealed preference rule* imposes only a subset of revealed preference constraints. At most one revealed preference constraint applies to each bid. In the clock stage, the bidder can shift to any package of the same or smaller size without constraint. However, the bidder can only shift to larger packages that satisfy revealed preference with respect to the prior bid; that is, at time t, if package q_t is larger than package q_{t-1}, then q_t has become relatively cheaper than q_{t-1}:

(P′) $$q_{t'} \cdot (p_{t'} - p_t) \leq q_t \cdot (p_{t'} - p_t).$$

Moreover, all supplementary bids $b(q)$ are capped by a single revealed preference constraint. Packages q the same size or smaller than q_f are capped by revealed preference with respect to the final clock package q_f:

(S′) $$b(q) \leq b(q_f) + (q - q_f) \cdot p_f.$$

Packages q larger than q_f also must satisfy revealed preference with respect to the smaller package q_s (where round s is the first round a package smaller than q was bid for):

(S´´) $$b(q) \leq b(q_s) + (q - q_s) \cdot p_s.$$

Notice that all the constraints are revealed preference constraints, so this indeed is a simplified revealed preference rule. Also, notice that for each package, there is at most a single revealed preference constraint that limits the bid. Thus, conforming to the rule is a simple one-dimensional problem. Just as with the eligibility point rule, there is no need for optimization tools to help the bidder conform to the rule.

One of the desirable features of the rule is that the final package in the clock stage plays an especially important role in limiting bids on all packages that are not larger than the final package. Thus, any distortion from profit maximization in the final clock package is especially costly to the bidder. Of course, the bidder never knows, which clock round will be the last, so there is always some incentive to bid consistent with profit maximization. Moreover, as excess demand falls, the probability that the current round will be the last tends to increase, strengthening the incentive for straightforward bidding throughout the clock stage.

Figure 14. Comparison of theory, simulation, and experiment in the 2.6 GHz auction

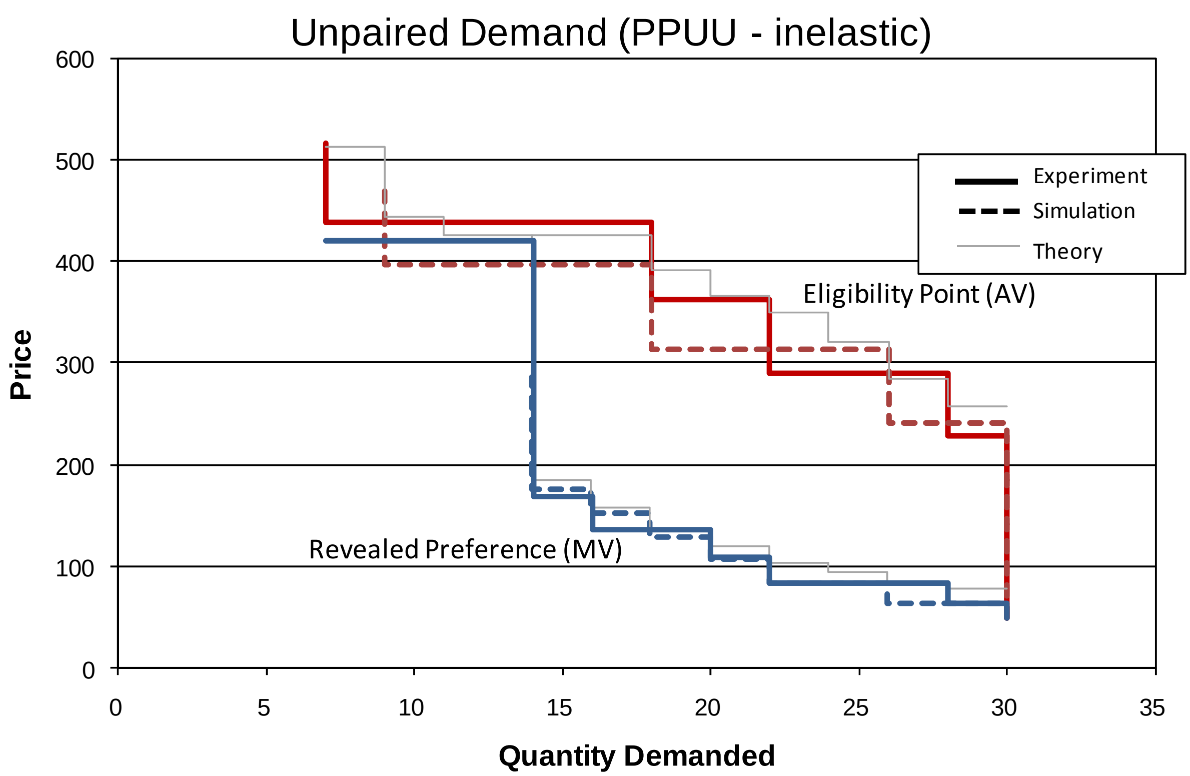

I tested a version of the simplified revealed preference rule for the UK 2.6 GHz auction. The methodology was as in the earlier tests. First I developed a number of realistic scenarios. The two activity rules were then evaluated using a theoretical clock auction with continuous prices, a simulation with discrete increments, and then in the experimental lab with the same experienced PhD students that participated in the earlier tests. Figure 14 compares the demand curves as bid in the clock stage under the two activity rules for a typical scenario. The subjects bid in a manner consistent with the assumed behavior in both the theory and the simulation. From this I conclude that the simplified revealed preference rule has all of the desirable properties of the revealed preference rule, without the complexity.

Ofcom has adopted a version of the simplified revealed preference rule for the 2.6 GHz auction.

7 Conclusion

The package clock auction is a large advance over the standard simultaneous ascending auction. It eliminates the exposure problem, it eliminates most gaming behavior, it enhances substitution, and it encourages competition. Most importantly, the package clock auction enables a technology neutral auction, which should be especially important with respect to the digital dividend. The auction, through the competitive bids, determines how the spectrum is organized, rather than the regulator. In an environment where the regulator has only a rumor of an idea about what technology or use is best, letting the auction resolve such matters can greatly expand the realized value of the scarce spectrum resource.

A further advantage of the package clock auction is that it is readily customized for a variety of settings. Typically, a communications regulator will have a sequence of auctions over many years, as new spectrum gradually is made available. The package clock auction can be adapted to the unique characteristics of any particular auction. Adopting a consistent and flexible auction platform reduces transaction costs for the government and, more importantly, the bidders.

The auction design also enhances competition. The process is highly transparent and encourages price discovery. There is enhanced substitution both through the product design and the auction format. Bidder participation costs are reduced.

As in any market design problem, an important task for the regulator is to identify and mitigate potential market failures. In this setting and many others, the most important potential failure is market power. This is especially an issue in settings where there already is a highly concentrated communications market and the spectrum is an essential input for any new entrant. The approach here allows the regulator to address this potential market failure, as well as others, with a variety of instruments, such as spectrum caps, set asides, or bidding credits. The instruments must be used with care, or else they may do more harm than good.

One of the greatest harms is delaying the allocation and award of spectrum. Avoiding economic loss from delay should be a main priority of the regulator. Incumbents often will argue that spectrum awards should be put off. Such arguments may simply be a far less costly means of impeding competition than outbidding an entrant in an auction.

Fortunately, the use of a state-of-the-art auction design, such as the package clock auction and its variants, does not cause delay. These auctions can be designed and implemented, even by developing countries, in short order, provided the country is using successful techniques adopted elsewhere. The bottleneck is regulatory procedures, not auction design and implementation. Providers of auction services can readily meet deadlines of a few months if necessary.

The package clock auction can be applied in many other industries. For example, the approach was proposed and tested for the auctioning of takeoff and landing slots at New York City's airports. The approach is well-suited for any setting in which there are many interrelated items, some of which are substitutes and some of which are complements.

More broadly, the approach described here is an example of using auction design to harness the power of markets. The approach leads to improved pricing of a scarce resource, improving decision making, both short term and long term. Innovation is fostered from the better pricing and assignment of the scarce resource.

References

Ausubel, Lawrence M. (2004), "An Efficient Ascending-Bid Auction for Multiple Objects," *American Economic Review*, 94:5, 1452-1475.

Ausubel, Lawrence M. (2006), "An Efficient Dynamic Auction for Heterogeneous Commodities," *American Economic Review*, 96:3, 602-629.

Ausubel, Lawrence M. and Peter Cramton (2002), "Demand Reduction and Inefficiency in Multi-Unit Auctions," University of Maryland Working Paper 9607, revised July 2002.

Ausubel, Lawrence M. and Peter Cramton (2004), "Auctioning Many Divisible Goods," *Journal of the European Economic Association*, 2, 480-493, April-May.

Ausubel, Lawrence M., Peter Cramton, R. Preston McAfee, and John McMillan (1997), "Synergies in Wireless Telephony: Evidence from the Broadband PCS Auctions," *Journal of Economics and Management Strategy*, 6:3, 497-527.

Ausubel, Lawrence M., Peter Cramton, and Paul Milgrom (2006), "The Clock-Proxy Auction: A Practical Combinatorial Auction Design," in Peter Cramton, Yoav Shoham, and Richard Steinberg (eds.), *Combinatorial Auctions*, Chapter 5, 115-138, MIT Press.

Ausubel, Lawrence M. and Paul Milgrom (2002), "Ascending Auctions with Package Bidding," *Frontiers of Theoretical Economics*, 1: 1–45, www.bepress.com/bejte/frontiers/vol1/iss1/art1.

Ausubel, Lawrence M. and Paul Milgrom (2006a), "Ascending Proxy Auctions," in Peter Cramton, Yoav Shoham, and Richard Steinberg (eds.), *Combinatorial Auctions*, Chapter 3, 79-98, MIT Press.

Ausubel, Lawrence M. and Paul Milgrom (2006b), "The Lovely but Lonely Vickrey Auction," in Peter Cramton, Yoav Shoham, and Richard Steinberg (eds.), *Combinatorial Auctions*, Chapter 1, 17-40, MIT Press.

Brusco, Sandro and Giuseppe Lopomo (2002), "Collusion via Signalling in Simultaneous Ascending Bid Auctions with Heterogeneous Objects, with and without Complementarities," *Review of Economic Studies*, 69, 407-436.

Bulow, Jeremy, Jonathan Levin, and Paul Milgrom (2009), "Winning Play in Spectrum Auctions," Working Paper, Stanford University.

Coase, Ronald H. (1959), "The Federal Communications Commission," *Journal of Law and Economics*, 2, 1-40.

Cramton, Peter (1995), "Money Out of Thin Air: The Nationwide Narrowband PCS Auction," *Journal of Economics and Management Strategy*, 4, 267-343.

Cramton, Peter (1997), "The FCC Spectrum Auctions: An Early Assessment," *Journal of Economics and Management Strategy*, 6:3, 431-495.

Cramton, Peter (2006), "Simultaneous Ascending Auctions," in Peter Cramton, Yoav Shoham, and Richard Steinberg (eds.), *Combinatorial Auctions*, Chapter 4, 99-114, MIT Press.

Cramton, Peter, Evan Kwerel, and John Williams (1998), "Efficient Relocation of Spectrum Incumbents," *Journal of Law and Economics*, 41, 647-675.

Cramton, Peter and Jesse Schwartz (2002), "Collusive Bidding in the FCC Spectrum Auctions," *Contributions to Economic Analysis & Policy*, 1:1, www.bepress.com/bejeap/contributions/vol1/iss1/art11.

Cramton, Peter, Yoav Shoham, and Richard Steinberg (2006), *Combinatorial Auctions*, Cambridge, MA: MIT Press.

Cramton, Peter, Andrzej Skrzypacz and Robert Wilson (2007), "The 700 MHz Spectrum Auction: An Opportunity to Protect Competition In a Consolidating Industry" submitted to the U.S. Department of Justice, Antitrust Division.

Day, Robert and Peter Cramton (2008), "The Quadratic Core-Selecting Payment Rule for Combinatorial Auctions," Working Paper, University of Maryland.

Day, Robert and Paul Milgrom (2008), "Core-selecting Package Auctions," *International Journal of Game Theory*, 36, 393-407.

Day, Robert W. and S. Raghavan (2007), "Fair Payments for Efficient Allocations in Public Sector Combinatorial Auctions," *Management Science*, 53, 1389-1406.

Federal Communications Commission (2002), "Spectrum Policy Task Force," ET Docket No. 02-135.

Harsha, Pavithra, Cynthia Barnhart, David C. Parkes, Haoqi Zhang (2009), "Strong Activity Rules for Iterative Combinatorial Auctions," *Management Science*, forthcoming.

Klemperer, Paul (2004), *Auctions: Theory and Practice*, Princeton University Press.

Kwasnica, Anthony M., John O. Ledyard, Dave Porter, and Christine DeMartini (2005), "A New and Improved Design for Multi-Object Iterative Auctions," *Management Science*, 51: 419 - 434.

McMillan, John (1994), "Selling Spectrum Rights," *Journal of Economic Perspectives*, 8, 145-162.

Milgrom, Paul (2004), *Putting Auction Theory to Work*, Cambridge: Cambridge University Press.

Milgrom, Paul (2007), "Package Auctions and Exchanges," *Econometrica*, 75, 935-966.

Milgrom, Paul (2009), "Simplified Mechanisms with Applications to Sponsored Search and Package Auctions," *Games and Economic Behavior*, forthcoming.

Pagnozzi, Marco (2007), "Should Speculators Be Welcomed in Auctions?" Working Paper, Universita di Napoli Federico II.

Parkes, David C. (2006), "Iterative Combinatorial Auctions," in Peter Cramton, Yoav Shoham, and Richard Steinberg (eds.), *Combinatorial Auctions*, Chapter 2, 41-78, MIT Press.

Porter, David, Stephen Rassenti, Anil Roopnarine, and Vernon Smith (2003), "Combinatorial Auction Design," *Preceedings of the National Academy of Sciences*, 100, 11153-11157.