A Practical Guide to the Combinatorial Clock Auction[*]

Lawrence M. Ausubel and Oleg Baranov

The Combinatorial Clock Auction (CCA) is an important recent innovation in auction design that has been utilised for many spectrum auctions worldwide. While the theoretical foundations of the CCA are described in a growing literature, many of the practical implementation choices are neglected. In this article, we examine some of the most critical practical decisions for a regulator implementing the CCA. Topics include: implementation of reserve prices; endogenous band plans; supplementary-round activity rules; competition policy; bidding languages; and allocation of the core burden. We illustrate our discussion with examples from recent spectrum auctions that used the CCA format.

Keywords: Auction Design; Combinatorial Clock Auction; Spectrum Auctions

JEL classification codes: D44; D47

---

[*] Corresponding author: Oleg Baranov, Department of Economics, University of Colorado, Boulder, CO, 80309, USA. Email: oleg.baranov@colorado.edu

We thank the editor, anonymous referees, Thomas Greve, Maarten Janssen and Sascha Schweitzer for helpful comments and suggestions. All errors are our own.

Since its proposal in a 2006 academic paper,[1] the combinatorial clock auction (CCA) has rapidly established itself as one of the leading formats for government auctions of telecommunications spectrum. Its initial implementations were for relatively small auctions and some of these applications may be viewed as experimental. However, in the past few years, usage of the CCA has gained substantial momentum. From 2012 to this writing in 2015, the CCA has been used for more than ten major spectrum auctions worldwide, allocating prime sub-1-GHz spectrum on three continents and raising approximately $20 billion in revenues (see Table 1). Despite the presence of an existing auction format—the simultaneous multiple round auction (SMRA)—which often performs reasonably well, the CCA has the potential of displacing it and becoming the new standard design choice for spectrum auctions.

Table 1: Combinatorial Clock Auctions to date, as of 2016

Country and Auction	**Year**	**Revenues**
Trinidad and Tobago Spectrum Auction	2005	$25.1 million ($US)
UK 10 – 40 GHz Auction	2008	£1.43 million
UK L-Band Auction	2008	£8.33 million
Netherlands 2.6 GHz Spectrum Auction	2010	€2.63 million
Denmark 2.5 GHz Spectrum Auction	2010	DKK 1.01 billion
Austria 2.6 GHz Spectrum Auction	2010	€39.5 million
Switzerland Spectrum Auction	2012	CHF 996 million
Denmark 800 MHz Spectrum Auction	2012	DKK 739 million
Ireland Multi-Band Spectrum Auction	2012	€482 million
Netherlands Multi-Band Spectrum Auction	2012	€3.80 billion
UK 4G Spectrum Auction	2013	£2.34 billion
Australia Digital Dividend Spectrum Auction	2013	$1.96 billion ($AU)
Austria Multi-Band Spectrum Auction	2013	€2.01 billion
Slovakia 800, 1800 and 2600 MHz Spectrum Auction	2013	€164 million
Canada 700 MHz Spectrum Auction	2014	$5.27 billion ($CA)
Slovenia Multi-Band Spectrum Auction	2014	€149 million
Canada 2500 MHz Spectrum Auction	2015	$755 million ($CA)

---

[1] The CCA format was proposed by Ausubel *et al*. (2006), first presented at the FCC's Conference on Combinatorial Bidding in Wye River, Maryland in November 2003. See: http://wireless.fcc.gov/auctions/default.htm?job=conference_presentations&y=2003 (last accessed: 30 March 2016).

The CCA design consists of a two-stage bidding process. The first stage, known as the ***clock rounds***, is a multiple-round clock auction. In each round, the auctioneer announces prices for all items and bidders respond with quantities demanded at these prices. If aggregate demand exceeds available supply for any items, the auctioneer announces higher prices for these items in the next round. The bidding process continues until prices reach a level at which aggregate demand is less than or equal to supply for every item. The second stage, known as the ***supplementary round***, is a sealed-bid auction process in which bidders can improve their bids made in the first stage and submit additional bids as desired for other combinations of items. Throughout the entire auction, all bids are treated as all-or-nothing package bids.

To determine winnings and associated payments, all bids placed during the clock rounds and all bids placed in the supplementary round are entered together into a standard winner determination problem (WDP). Winning packages are determined by finding an allocation that maximizes the total value (as reflected in bids) subject to feasibility constraints: each item can be sold only once and only one bid from each bidder can be selected as part of the winning allocation. Corresponding payments are found by solving a series of counterfactual WDPs that identify the relevant opportunity costs (second prices) imposed by winners on other participants. In general, a payment rule based on opportunity costs creates incentives for bidders to reveal their true values, facilitating efficient outcomes.

In current practice, it is standard for the CCA also to include a third bidding stage, known as the ***assignment stage***. This stage is added to the CCA in order to significantly simplify bidding in the first two stages. The main idea is to treat several closely-related items as completely identical during the clock and supplementary rounds. For example, the European digital dividend auctions auctioned six distinct licenses in the 800 MHz band. In these auctions, bidders were typically asked to bid for quantities of “generic” 800 MHz spectrum blocks during the first two stages of the auction. In a third stage that takes the winning allocations of the “generic” spectrum as given, bidders had the opportunity to compete for specific frequency assignments within the 800 MHz band.

The assignment stage, usually implemented as a sealed-bid auction, determines the mapping from generic spectrum to physical frequencies.[2]

The literature on the CCA is growing rapidly. Ausubel and Baranov (2014) describe the evolution of the CCA, including expected innovations to the design. Cramton (2013) outlines the flaws in the SMRA design and argues how the CCA design solves them. Various pricing mechanisms for combinatorial auctions and their properties, including the pricing rules currently used for CCAs, have been studied by Parkes (2001), Ausubel and Milgrom (2002), Day and Raghavan (2007), Day and Milgrom (2008), and Day and Cramton (2012). Possibilities for strategic manipulations of the CCA are explored by Janssen and Karamychev (2013) and Levin and Skrzypacz (2014). Experimental comparisons of the CCA design with other auction designs have been studied by Bichler *et al*. (2013) and Bichler *et al*. (2014).

Despite this burgeoning literature, there has been little written to address some of the most pressing practical issues facing regulators. The goal of the current article is to fill some of this void. Section 1 addresses reserve prices. In the consultation processes that frequently precede spectrum auctions, a great deal of attention is dedicated to the *levels* of reserve prices. However, very little thought typically goes into the *implementation* of reserve prices, which in CCAs can be as consequential as the levels. As a result, many CCAs have adopted implementations that unnecessarily distort bidders' incentives. Section 2 discusses endogenous band plans. In current practice, the configuration of the spectrum is largely determined through lobbying by stakeholders in a consultation process preceding the auction. One of the promises of combinatorial auction formats such as the CCA is that they open the possibility of the "market" not only determining the allocation of spectrum but also the underlying band plan. We point out some issues and potential pitfalls in designing a CCA with endogenous band plans. Section 3 examines activity rules in the supplementary round. Activity rules make the early rounds of dynamic auctions informative, while still allowing bidders to switch to bidding on different items in later rounds. In a CCA, they serve the additional role of limiting overbidding and surprise bids in the supplementary

---

[2] For more detailed descriptions of the CCA mechanics, see regulator websites such as: https://www.ic.gc.ca/eic/site/smt-gst.nsf/eng/sf10583.html and http://www.acma.gov.au/Industry/Spectrum/Digital-Dividend-700MHz-and-25Gz-Auction/Reallocation/combinatorial-clock-auctions-reallocation-acma (last accessed: 30 March 2016).

round; we provide some guidance to currently used activity rules. Section 4 considers competition policy. When competition policies (or any other side objectives) are integrated into the auction, important properties of the CCA design need to be preserved. We show by example that some recent European CCAs appear to have had incomplete integrations. Section 5 treats briefly the two topics of bidding language and the allocation of the "core burden". While bidding language may appear to be an abstract issue, regulators implementing CCAs have always imposed a maximum number of bids that can be submitted in the supplementary round. Bidding languages provide an opportunity to make the bid limit non-binding. The allocation of the difference between Vickrey-Clark-Groves and core prices may also appear overly academic, but it affects the fundamental fairness between large and small bidders, and it can be improved with a simple fix involving weighting. Section 6 concludes.

1 Reserve Prices

Reserve prices have been employed in the vast majority of spectrum auctions, irrespective of their formats. At least three rationales have been put forward for the use of reserve prices. First, they establish a lower bound for auction revenues. Second, they promote competitive behavior among bidders by reducing possible gains from various strategic manipulations. In public auctions, there is a third rationale: reserve prices internalise the societal opportunity cost of selling assets to private parties now rather than retaining them for the future.

In the consultation processes that frequently precede spectrum auctions, a great deal of attention is dedicated to the *levels* of reserve prices. However, comparatively little thought typically goes into the *implementation* of reserve prices, which in CCAs can be as consequential as the levels.[3] The explanation for the importance of the reserve-price implementation in CCAs is quite straightforward. In previous formats used for spectrum auctions

---

[3] For example, in the run-up to the UK 4G Auction, Ofcom (the UK regulator) devoted 85 paragraphs (¶8.01 – ¶8.85) of its final major consultation document to the determination of levels of reserve prices, but only two paragraphs (¶7.35 – ¶7.36) to its change in implementation of reserve prices from the first alternative described below to the second. See: http://stakeholders.ofcom.org.uk/binaries/consultations/award-800mhz/statement/statement.pdf (last accessed: 30 March 2016).

(e.g. the SMRA), there are individual prices paid for each item sold, and so starting at the reserve price for each item has an unambiguous meaning. By contrast, auctions with package bidding (e.g. the CCA) result in prices for packages of items, which do not have linear representations as sums of single-item prices.

In this section, we focus on the two approaches to reserve prices that have been used in CCAs:

> ***Bundle reserve prices***: This is the simple requirement that the payment for any bundle of goods must be at least the sum of the reserve prices established for the items contained in the bundle.
>
> ***Incremental reserve prices***: Under this approach, reserve prices are interpreted as minimum incremental costs of acquiring additional items. Incremental reserve prices ensure that an increase in payment for winning extra items is always at least the reserve prices for these items.

We illustrate the difference using a simple example with two items, A and B, each with a reserve price of 10. Suppose that, absent the reserve prices, a bidder would pay 15 for item A, 14 for item B, and 19 for A+B (the combination). Given the reserve prices, a payment of 15 for item A or a payment of 14 for item B is adequate, while a payment of 19 for A+B is too low. Under the bundle reserve prices, the payment for A+B is increased to 20, the sum of the reserve prices for the two items. Under the incremental reserve prices, the payment for A+B is increased to 25, increasing the stand-alone payment of 15 for item A by the reserve price of item B. Following directly from definitions, incremental reserve prices always satisfy the bundle reserve prices, but the reverse implication does not hold.

To implement bundle reserve prices, the regulator might employ an approach known as “bounds only” — the submitted bid for any package must be at least the sum of its component reserve price and, if the opportunity-cost-based price for a bidder ends up being too small, the final payment for any winning package is increased to its component reserve prices. To implement incremental reserve prices, the regulator might use an approach known as “reserve bidders” — fictitious bidders who bid for each individual item at its reserve price are added to the winner determination and price determination processes, thus explicitly applying opportunity costs at the reserve price level to all possible combinations of items.

It might seem that the incremental approach would necessarily generate at least as much revenue as the bundle approach, but Day and Cramton (2012) show that there is no general revenue ranking of the two treatments. The incremental approach would indeed generate at least as much revenue as the bundle approach when they both allocate the same set of items to actual bidders. However, the incremental approach may generate lower revenues if it allocates fewer items.

There are two clear benefits of bundle reserve prices for a regulator. First, the bundle approach always results in a weakly larger set of items being sold, and the number of unallocated items is frequently viewed by regulators as a measure of failure. Second, final payments by bidders are less sensitive to the particular choice of reserve prices under the bundle approach. Or putting it differently, reserve prices serve only to protect aggregate revenues — once opportunity-cost prices have become sufficiently high, the exact choice of reserve prices has no further effect on bidders' payments. By contrast, incremental reserve prices might affect bidder payments even with reasonable competition for some of the items. Given that many regulators set reserve prices using nontransparent ad-hoc procedures, a less intrusive implementation is viewed as more desirable.

At the same time, the incremental approach is preferred from the perspective of bidder incentives. Use of opportunity-cost-based pricing in the CCA incentivises bidders to bid their true values.[4] It turns out that these incentives can be damaged by bundle reserve prices as implemented via the "bounds only" approach.

Consider our earlier example with items A and B. Suppose that the values of Bidder 1 are $v_1(A) = 15, v_1(B) = 20, v_1(AB) = 25$ and that the values for Bidder 2 are $v_2(A) = 15, v_2(B) = 19, v_2(AB) = 25$. If both bidders bid truthfully, Bidder 2 is awarded item A and her opportunity cost payment is 5 (the value taken from Bidder 1 by Bidder 2's participation).[5] If each item has a reserve price of 10, under the bundle approach, Bidder 2's payment will be increased to 10 to meet the reserve price. Now suppose that Bidder 2 deviates and bids 21

---

[4] CCAs traditionally use core-selecting payment rules that might differ from the Vickrey payment rule in specific scenarios. When this is the case, bidders sometimes have incentives to understate their values for some bundles.

[5] If Bidder 2 does not participate, Bidder 1 would have won AB instead of B for additional value of 5 $\left(v_1(AB) - v_1(B) = 5\right)$.

for item B, resulting in an inefficient allocation. Then Bidder 2 will be awarded a more valuable item B while still paying only 10 (the opportunity cost of winning item B). By contrast, the same deviation by Bidder 2 under the incremental approach will result in a payment of 15, rendering the deviation unprofitable. The last observation is general: the incremental approach never disturbs the incentives to bid truthfully and, therefore, is more likely to lead to efficient outcomes.

Intuitively, bundle reserve prices, when binding, generate a "reserve credit" that can be used by bidders to purchase extra items or to exchange current items for more valuable ones. In the previous paragraph, when winning item A, Bidder 2 has a reserve credit of 5, and the incremental cost of getting item B instead of A is also 5. Hence, Bidder 2 can switch from winning item A to winning item B for free by inflating its bid for B.

Generally, overstating values can be risky when opponents' values are private information. In sealed-bid auctions, such uncertainty can mitigate the incentives to overstate values. However, in the CCA, bidders are able to make inferences about opponent's values from the aggregate demand information they get during the clock rounds. For example, if at opening clock prices of (10, 10), Bidder 1 demands item B only, Bidder 2 would infer that Bidder 1's incremental value for A is less than 10, which in turn makes overbidding on item B a relatively safe bet.

The incentives for overbidding can be rather large, as they appear to have been in the 2013 Slovakian 4G spectrum auction. In the 1800 MHz band, three large "B" blocks were set aside for an entrant. In the prime 800 MHz band, the supply of "A" blocks was six, and the three incumbents were subject to spectrum caps of two blocks each. In the less valuable 2.6 GHz band, the supply of "C" blocks was 14, and there was no spectrum cap. The reserve price for an "A" block was €19 million and the reserve price for a "C" block was €1 million.[6]

Observe that, in such a scenario, if only incumbents competed for the "A" blocks, there would be zero opportunity cost associated with winning "A" blocks (since the aggregate demand would never exceed the

---

[6] See the published auction rules at http://www.teleoff.gov.sk/data/files/35571.pdf and http://www.teleoff.gov.sk/data/files/33771.pdf (last accessed: 30 March 2016).

supply). This implies a €38 million "reserve credit" for each incumbent that could be potentially used for rampant overbidding on "C" blocks.

Examples of CCAs that adopted a bundle reserve approach were in the UK (two auctions held in 2008), Denmark (2010 and 2012), Switzerland (2012), Ireland (2012) and Slovakia (2013). An incremental reserve price implementation was chosen for CCAs in UK 4G (2013),[7] Australia (2013) and Canada (2014 and 2015).

2 Endogenous Band Plans

Radio spectrum can be used for multiple purposes. In current practice, the configuration and permitted use of the spectrum is largely determined through lobbying by stakeholders in a consultation process preceding the auction. One of the promises of combinatorial auction formats such as the CCA is that they open the possibility of the "market" not only determining the allocation of spectrum but also the underlying band plan. A regulator who wishes to accommodate bidders who plan to use spectrum in mutually-exclusive ways can consider making the band plan endogenous. This approach allows the same physical spectrum to be offered multiple times through mutually-exclusive lots, making the auction attractive for bidders with diverse interests. The CCA design makes it very easy for regulators to embed endogenous band plans into their auctions.

A classic example of an endogenous band plan embedded within a CCA is the initial design for the 2.6 GHz UK auction in 2007.[8] The proposal included 38 blocks (5 MHz each) in the 2.6 GHz band, with endogenous determination of the number of paired lots (spectrum blocks suitable for frequency division duplex or FDD) and the number of unpaired lots (blocks suitable for time division duplex or TDD). Note that FDD lots and TDD lots need to be physically separated from each other, so once the band plan is determined, the quantities of each are locked in. According to the auction rules, bidders would submit bids specifying both the number of blocks they

---

[7] Ofcom initially proposed to use the bundle approach for all lots, but later decided against it — see footnote 3. The incremental approach was applied to the A1, A2, C and E lots. The bundle approach was applied to the D1 and D2 lots.

[8] This design was never implemented as the auction was later superseded by the 2013 UK 4G Auction that additionally included the 800 MHz and 1800 MHz bands. The expansion of the supply coupled with a significant time delay warranted a complete redesign of the auction.

wanted and whether they wanted them in paired, unpaired, or combined configurations. The auctioneer would then determine the actual split between the two technologies by maximizing the total value of the allocation.

A more exotic example of an endogenous band plan was included in the actual UK 4G Auction held in 2013. There, up to four of the paired lots in the 2.6 GHz band were made available either for low-power shared use (as D1 or D2 lots) or as high-power non-shared use (as C lots). In the low-power use, up to 10 different operators would share the same frequencies. In the high-power use, each C lot would be owned and operated by a single operator. The final decision between low-power and high-power usage was to be determined by comparing the sum of bids from bidders desiring low-power lots with the bids from bidders desiring high-power lots.

While endogenous band plans are conceptually attractive, a regulator should be aware of three types of complications that they may create.

Interaction with Reserve Prices

When a regulator incorporates extra design elements such as endogenous band plans, the interaction of the reserve price policy with the new design elements requires careful reexamination. For example, Ofcom in its UK 4G Auction decided to use the "reserve bidders" approach for all lots except the lots designated for the low-power shared use (D1 and D2). Instead, the low-power lots were subject to the "bounds only" approach. The rationale behind Ofcom's decision appears to be reasonable. On the one hand, Ofcom avoided situations in which the fictitious reserve bidders for D lots would have helped actual bidders bidding on D lots to compete against the C lot bidders. On the other hand, Ofcom's policy introduced a positive bias into its objective of "technological neutrality": fictitious "reserve bidders" for C blocks can displace actual bidders for D lots even if the latter bid above their reserve prices and should be awarded their spectrum.

Free-Rider Problem

All CCAs to date used core-selecting payment rules that are well known to induce certain incentives for free riding amongst bidders. In recent CCAs for spectrum, these incentives were never a major concern.[9] However,

---

[9] The typical market structure of the wireless industry, with 3 to 4 major incumbents, combined with common spectrum caps almost always guaranteed that bidders would not benefit from such free-riding strategies. The only potential

the free-rider problem can be quite severe when the auctioneer uses a combination of a core-selecting payment rule and an endogenous band plan. If one set of bidders bids on licenses for use with technology A and another set of bidders bids on the same frequencies for use with technology B, bidders within each group can have strong incentives to limit their bidding and to free ride on other bidders seeking to use the same technology.

The free-rider problem in the "FDD vs. TDD" example above appears to be minor, but the "Low-Power vs. High-Power" example looks problematic. Up to 10 bidders for low-power use were effectively invited to compete jointly against one or two bidders for high-power use, inducing a severe free-rider problem among the low-power bidders. All would have strong incentives to bid very conservatively and to let other low-power bidders pick up the tab. Therefore, from the planning stages of this auction, the low-power shared technology was unlikely to prevail. In fact, the bidding data from the auction shows that one bidder, Niche (BT), placed a substantial number of supplementary bids that were suggestive of extreme incentives for free riding.[10]

Strategic Manipulations

Another important issue that arises in the context of endogenous band plans is unintended opportunities for strategic manipulations. In many circumstances, a bidder in a CCA may find it beneficial to manipulate the price trajectory during the clock rounds in order to put itself in a better position for the supplementary round.[11]

The UK 4G Auction was ripe for such strategies. In some circumstances, a bid for a single low-power D2 block — in combination with the price incrementing policy that was used in the auction — would reduce the endogenous supply of high-power C blocks and cause an artificial increase in their clock price.

---

exceptions to date may have been auctions with regional licenses in countries with strong regional bidders, which naturally may create a free-rider problem.

[10] Niche placed 54 (out of 89) bids for packages that included either a D1 or D2 lot, that were just a minimal increment (£1000 or £2000) over its corresponding base bids without D lots. The amounts of the base bids were at least £20 million.

[11] Clock prices affect revealed-preference constraints that in turn set upper bounds for the amounts of supplementary bids (see Section 3 for details).

3 Supplementary-Round Activity Rules

The imposition of activity rules in dynamic auctions has been one of the most important innovations in recent auction design. Generally speaking, activity rules are intended both to speed up the auction process and to curtail "bid sniping" opportunities (bidders concealing their true intentions until the very end of the auction). The traditional notion of "bid sniping" comes from eBay auctions with fixed ending times, where the high-value bidder can sometimes reduce its payment by submitting its winning bid at the very last moment. In a CCA, activity rules serve the additional roles of limiting surprise bids in the supplementary round and the bidder's ability to drive up its opponents' payments by overbidding for packages that can no longer be won.

CCAs to date have had considerable diversity in their activity rules. There are two distinct places in the CCA where the regulator needs to select an activity rule: (1) a clock-round activity rule that limits the set of items on which the bidder can bid in later clock rounds, based on bids in earlier clock rounds; and (2) a supplementary-round activity rule that limits the amounts that the bidder can bid on various packages in the supplementary round. Clock-round activity rules (and issues thereof) are in essence a special case of general considerations relating to any dynamic auction — the next two paragraphs will provide an overview of the issues. However, supplementary-round activity rules are relatively unique to the CCA. Therefore, the main focus of this section will be to provide the regulator with some guidance on selecting the supplementary-round activity rule.

Historically, spectrum auctions have utilised points-based activity rules: the auctioneer assigns a number of points to each item, before the start of the auction, and requires each bidder to adhere to monotonicity in the total number of points associated with each successive bid. In other words, a bidder must bid for a large bundle in early rounds in order to be allowed to bid for an equally-large bundle in later rounds, limiting possibilities for bid-sniping. Unfortunately, such an approach may also interfere with truthful bidding. For this reason, Ausubel *et al*. (2006) suggested introducing revealed-preference considerations into the clock-round activity rule of the CCA. Early implementations of the CCA tended to ignore their suggestion and to follow the traditional points-based approach in clock rounds. More recently, a hybrid approach, based upon a combination of points and revealed-preference considerations, has been taken in several spectrum auctions.

In Ausubel and Baranov (2014; 2016a), we show that a hybrid approach can actually make matters worse, and we propose instead to base the clock-round activity rule entirely upon the Generalized Axiom of Revealed Preference (GARP) — a well-known rationality concept used in economics.[12] Furthermore, in Ausubel and Baranov (2016b), we show that the GARP-based activity rule can be a foundation for a pricing mechanism that dynamically approximates VCG payoffs and thereby improves bidding incentives in the CCA.

For the supplementary round, the CCA's activity rule is also based on a combination of points and revealed-preference ideas. Current implementations of activity rules rely on the concepts of the Relative Cap, the Intermediate Cap and the Final Cap. All of them use the following concept of revealed-preference constraint:

> ***Revealed-Preference Constraint:*** The revealed preference constraint for package x with respect to the clock round t is $b(x) \leq b(x_t) + p_t(x - x_t)$, where $b(x)$ is the bid amount for package x, x_t is the package demanded in Round t, $b(x_t)$ is the final bid amount for package x_t, and p_t is the vector of clock prices for Round t.[13]

Intuitively, the revealed-preference constraint states that it is unreasonable for the bidder to claim a high value for package x relative to package x_t, given that the package x_t was revealed to be preferred to package x in Round t. The Relative Cap, Intermediate Cap and Final Cap rules are defined in terms of revealed-preference constraints with respect to certain clock rounds as follows:

---

[12] GARP imposes revealed-preference constraints against the bidder's demands in all clock rounds. To put it differently, the GARP activity rule requires the bidder to exhibit rational behavior in all of its demand choices. One significant advantage of the GARP activity rule to regulators is that it completely eliminates the need for points — and therefore it completely eliminates the need to assign points. Furthermore, the GARP activity rule does not require a monotonic price trajectory.

[13] More precisely, this is the constraint imposed by the Weak Axiom of Revealed Preference (WARP) for a bidder with a quasilinear payoff function.

> ***Relative Cap:*** A bid for the package *x* should satisfy the revealed-preference constraint with respect to the last clock round in which the bidder's eligibility, as measured by points, was at least the total points associated with package *x*.
>
> ***Intermediate Cap:*** A bid for the package *x* should satisfy the revealed-preference constraints with respect to all eligibility-reducing rounds starting from the last clock round in which the bidder's eligibility, as measured by points, was at least the total points associated with package *x*.
>
> ***Final Cap:*** A bid for the package *x* should satisfy the revealed-preference constraint with respect to the final clock round.

The Final Cap is the most natural constraint of revealed preference theory coming out of the clock rounds. Yet many regulators decide to rely only on the Relative Cap, and view the Final Cap as giving excessive allocation stability between the clock stage and the supplementary round. In the extreme case, it may be impossible for the supplementary bids to change the allocation of the final clock round; consequently, opportunity cost pricing based on these bids may become unreliable due to poor incentives. Recently, both the ACMA (the Australian regulator) and Ofcom (the UK regulator) used this rationale to decide against the Final Cap.

However, a simple omission of the Final Cap can be costly. While the absence of the Final Cap does mitigate the incentive issue, it does so at the expense of leaving substantial bid sniping opportunities and thus risking elimination of many advantages of a dynamic auction. The power of the Final Cap comes from its use of revealed-preference constraints generated *for all bidders at the same time* (more specifically, *using the same price vector*). In contrast, both the Relative Cap and Intermediate Cap use bidder-specific eligibility-reducing rounds to generate revealed-preference constraints.

To illustrate the importance of the Final Cap, we compare the extent of bid sniping opportunities available to bidders in the supplementary round under several activity rules. For the UK 4G Auction (see Table 2), we have calculated “theoretical”[14] bid amounts needed to protect final clock packages for all major bidders[15] under the Relative Cap (the actual activity rule used in the auction), the Intermediate Cap (considered, but never used), the Relative Cap + Final Cap (used in Ireland), and the Intermediate Cap + Final Cap (used in Canada). The exposure numbers were generally high in the UK 4G Auction, mostly due to the large value of unsold lots at the end of the clock stage. For this reason, we report two sets of numbers. “Exposure” provides the theoretical amounts required to protect the final clock package using the actual amounts of unsold lots in the auction, while “Net Exposure” provides the theoretical amounts when the effect of unsold lots on exposure is removed (as if the auctioneer had placed bids for all unallocated lots at the final clock prices). All exposure numbers are reported as percentages of the bidders’ final clock prices. As can be seen from the table, the Final Cap consistently reduces the amount that any bidder needs to bid to guarantee its winnings, due to reduced bid sniping opportunities. Furthermore, the reduction is rather significant. Even without undersell (Net Exposure), bidders might need to increase their bids by three-quarters or more in order to protect their final clock packages under the Relative Cap. As an extreme example, Niche would have needed to increase its bid by almost a factor of seven in order to assure that it won its final clock package. But the need to substantially increase bids after the clock rounds undermines the purpose of conducting the clock rounds at all. Our calculations demonstrate that the Relative and Intermediate Caps by themselves provide rather weak protection against bid sniping in the realistic setting of the actual UK 4G Auction data.

---

[14] These “theoretical” exposures are calculated as maximum incremental values that the bidder’s opponents could place on the lots in the bidder’s final clock package and unsold lots on top of values expressed for their final clock packages. We calculated them based on the full history of clock bids; while the disaggregated bids were unavailable to bidders during the auction, the calculation could be approximated using aggregate demand, which was disclosed. Also, since the low-power D1 and D2 lots did not enter into the final allocation, we excluded any effects associated with D1 and D2 lots from this calculation. Bid data for the UK 4G Auction used in this simulation is available at http://stakeholders.ofcom.org.uk/spectrum/spectrum-awards/awards-archive/completed-awards/800mhz-2.6ghz/auction-data/ (last accessed: 30 March 2016).

[15] They are calculated for all bidders who had non-empty final clock packages in the UK 4G Auction with the exception of Three. Under the competition policy employed in this auction, Three was guaranteed to win one of the designated minimum spectrum portfolios and therefore never needed to protect its bids for any of these packages (and was made aware of this fact during the auction).

Table 2: Exposure Calculation for the UK 4G Auction (2013)

Bidder / Final Clock Package	**Type**	**Relative Cap (used in UK)**	**Intermediate Cap**	**Relative Cap + Final Cap (used in Ireland)**	**Intermediate Cap + Final Cap (used in Canada)**
Vodafone	Exposure	245%	188%	170%	170%
(2 – A1, 3 – C)	Net Exposure	176%	118%	100%	100%
Telefonica[16]	Exposure	290%	216%	192%	192%
(1 – A2)	Net Exposure	198%	124%	100%	100%
EE	Exposure	715%	456%	456%	456%
(9 – E)	Net Exposure	359%	100%	100%	100%
Niche (BT)	Exposure	1103%	637%	524%	524%
(2 – C)	Net Exposure	679%	212%	100%	100%

Notes. All exposure numbers are reported as percentages of the bidders' final clock prices.

4 Competition Policy

One of the important tasks of a spectrum regulator is to incorporate competition policy into spectrum auctions. To be clear, in discussing "competition policy," we are generally not referring to introducing competition into the spectrum auctions themselves, but to injecting competition into the *downstream market* for mobile voice and data services. Spectrum is a scarce input into the provision of mobile services, and concentration in spectrum holdings would generally lead to concentration in the downstream market. To the extent that the spectrum auction facilitates new entry and reduces concentration in spectrum holdings, it can help to improve the competitive performance of the downstream market. As we will now see, the CCA is quite robust to the introduction of standard competition measures, and readily allows more flexible measures ("virtual set-asides") than are typically used today. However, the CCA becomes more sensitive when adding novel instruments — and it is also easy for the well-intentioned regulator to introduce inadvertent design flaws into the process.

Historically — and generally in conjunction with use of the SMRA format — the two most frequent instruments for competition policy in spectrum auctions have been the *set-aside* and the *spectrum cap*:

---

[16] In the actual UK 4G Auction, Telefonica was bidding for a (1–A2, 1–D2) package in rounds 41-52. For simplicity, we are omitting D1 and D2 lots from these calculations and assuming that Telefonica was bidding for (1–A2) lot in rounds 41-52. See also footnote 14.

Set-aside: One or more spectrum blocks are reserved for a particular class of bidders. The most common form that this has taken is that a specific block of spectrum is set aside for “entrants”; while “incumbents” are excluded from bidding on the set-aside block.[17]

Spectrum cap: A limit is placed on the quantity of spectrum within a given group of bands that a bidder is permitted to acquire or hold. Spectrum caps may be applied to winnings within a given auction or they may apply cumulatively to specified existing holdings as well as to acquisitions within the given auction. In countries such as the US, India, Canada and Australia, where spectrum licenses are regional, spectrum caps have been applied to each and every geographic region.[18]

However, due to the nature of older auction formats (e.g. the SMRA), there have been two frequent limitations on these policy instruments. First, since the bids have been for specific (as opposed to generic) licences, set-asides have needed to be for specific licences. This required regulators to make unnecessarily intrusive decisions (e.g. which exact frequencies should be reserved for entrants). Second, since the determination of high bids has historically been done at the individual license level, spectrum caps were applied individually to each incumbent (as opposed to establishing an aggregate cap on the overall acquisitions or holdings of all incumbents). This

---

[17] Set-asides have been used in auctions in many countries, including the US and the UK. In the US Broadband PCS spectrum auctions which began in 1994, the FCC set aside two of the original six blocks of broadband PCS spectrum for “designated entities” in an effort to promote small business ownership of spectrum. The implementation of set-asides did not work out well, as the set-aside was bundled with other policy instruments, such as instalment payments, that were intended to favour small businesses. The result of this exercise was that one of the largest winners of set-aside spectrum entered into bankruptcy, and there was a roughly ten-year period when most of the set-aside spectrum went unused. By contrast, the UK had an apparently successful experience with set-asides. In the UK 3G Auction of 2000, Ofcom set aside one of five 3G licenses for a new entrant. This resulted not only in actual entry and hence competition in the downstream market for wireless services, but in increasing competition in the auction and probably increasing auction revenues.

[18] Spectrum caps have been prevalent in spectrum auctions worldwide. At the time of the US Broadband PCS spectrum auctions, the FCC established a 45-MHz cap for commercial mobile radio spectrum, including both existing cellular licenses and the new PCS licenses being auctioned. By contrast, in many of the European 3G auctions in 2000, bidders were limited to winning at most a specified number of spectrum blocks, irrespective of their existing spectrum holdings. In recent 4G auctions, bidders have often been subject to an overall constraint on existing and new holdings for spectrum generally, together with a tighter constraint on the prime sub-1-GHz spectrum. For example, in the 2013 UK 4G spectrum auction, there was both a 105-MHz overall cap on existing and new holdings (which proved to be binding for incumbent Everything Everywhere) and a 27.5-MHz cap on sub-1-GHz spectrum (which proved to be binding for incumbents Vodafone and Telefonica).

potentially had the effect of impeding competition in the auction among incumbent bidders and thereby depressing auction revenues. Both of these preconditions change under the CCA.

Virtual Set-Asides/Aggregate Spectrum Caps

In new auction formats such as the CCA, it is possible to go beyond the two aforementioned limitations and to provide for a *virtual set-aside* (which may equivalently be viewed as an *aggregate spectrum cap*):

> ***Virtual set-aside/aggregate spectrum cap:*** If a quantity of N generic spectrum blocks is offered, then a particular class of bidders (e.g. "incumbents") is limited, in aggregate, to winning a quantity of $N - K$ of these generic blocks. Such a policy instrument can be called an *aggregate spectrum cap*. Equivalently, this can be viewed as reserving a quantity of K generic spectrum blocks for another class of bidders (e.g. "entrants"). Such a policy instrument can be called a *virtual set-aside*.

One key advantage of this policy instrument (as compared to the implicit spectrum reservation of a standard spectrum cap) is that, for any quantity reserved for entrants, there can now be greater competition among incumbents. For example, suppose that there are two incumbent mobile operators and six blocks of spectrum available. The regulator can implicitly reserve two blocks for entrants by limiting each incumbent to an in-auction spectrum cap of two blocks. However, unless entrants demand greater than two blocks, this instrument would eliminate all competition for the incumbents. Instead, the regulator can establish an aggregate cap of four blocks on the incumbents. Each one is free to bid for and win up to four blocks. With a virtual set-aside, there are still two blocks reserved for entrants, but competition between the two incumbents is permitted to occur.

A virtual set-aside in the CCA has several advantages over the standard set-aside in the SMRA. First, it is more flexible, and allows the market to determine which specific frequency blocks are won by entrants. Second, with the CCA's price determination mechanism rather than implicit uniform pricing, a single entrant may be capable of winning the entire set-aside, as opposed to being vulnerable to pressure from a second entrant who may force it to surrender some of the set-aside. (And the regulator may, for competition reasons, prefer that a single entrant win the entire set-aside, as a single substantial entrant is likely to pose a greater competitive threat to incumbents than a diffuse competitive fringe.) Third, in some situations, a standard set-aside may be viewed as an excessively strong policy and of questionable legality. While the aggregate spectrum cap is isomorphic to a virtual set-aside, the aggregate spectrum cap (as a limit on incumbents, rather than a reservation for entrants) may be considered more acceptable and less subject to legal challenge.

Spectrum Floor

The UK 4G Auction of 2013 went a step further beyond a virtual set-aside and established a "spectrum floor". Ofcom (the regulator) decided that either a 2x5 MHz block of 800 MHz spectrum or a 2x20 MHz block of 2.6 GHz spectrum would be reserved for an entrant. This went further than a virtual set-aside in two respects. First, it allowed the market to determine not only which frequency blocks would constitute the set-aside, but whether they would belong to the more valuable 800 MHz band or the less valuable 2.6 GHz band. Second, it institutionalized that the entire set-aside would be required to be won by a single entrant.

> ***Spectrum floor:*** Two or more alternative sets of spectrum ("minimum spectrum portfolios") are reserved for a particular class of bidders. The choice of which minimum spectrum portfolio is awarded, and to which bidder, is decided endogenously by the solution to a modified winner determination problem: maximize the value of accepted bids, subject to feasibility and to awarding one of the minimum spectrum portfolios to an eligible bidder.

With a CCA, two issues are likely to arise when spectrum is reserved for entrants. The first issue can be referred to as "infeasible bids". Given that certain spectrum blocks are reserved for entrants, certain other combinations of blocks might never be part of a feasible winning bid. In this event, placing bids for infeasible combinations of blocks might become a stalling tactic or a way to drive up prices of certain spectrum bands; therefore, it is essential that the auction rules prevent the submission of infeasible bids. The second issue relates to the supplementary bids. Under the rules generally used in CCAs today, the bidder is permitted to submit a supplementary bid that raises its final clock package (the package that it bid for in the final round of the clock stage) by any arbitrary amount. However, there is one exception: no supplementary bids are permitted to be submitted for the null set, even if it is the bidder's final clock package (i.e. if the bidder has dropped out of the clock stage by the final clock round). When spectrum is reserved for entrants, similar logic would suggest that the bidder should not be allowed to place a supplementary bid on the reserved package.[19]

5 Other Considerations

Compact Bidding Languages

The standard bidding language for CCAs treats each of a bidder's bids as an all-or-nothing package bid, with the restriction that only one bid per bidder can win. This bidding language is fully expressive, in that it allows bidders to communicate all possible valuations. However, it comes at a steep price. The bidding language is non-compact in that it may require a very large number of bids to express simple (e.g. additive) preferences.[20]

While the bidding language may appear to be an abstract issue, regulators implementing CCAs have always needed to impose a limit on the number of bids that each bidder can submit in the supplementary round.[21] Bidding languages provide an opportunity to make the bid limit non-binding.

One approach for auctions with large numbers of items offered, used successfully in Canada's 2500 MHz Auction (with 318 licences grouped into 106 categories), is to allow bidders to place non-mutually-exclusive "OR bids" as increments to their final clock packages. The list of must-address issues for the regulator then includes the interaction of OR bids with the CCA activity rules and the treatment of OR bids by the winner-determination problem.

---

[19] More generally, a clean solution to this problem might be to require the entire collection of highest bids to be consistent with a GARP-based activity rule (see the discussion in Section 3).

[20] See Nisan (2006) for a general overview of bidding languages in combinatorial auctions.

[21] The reason for this limit is to assure that the winner determination problem (which is NP-complete) can be solved in reasonable time. The UK 4G Auction set a limit of 4000 bids, while Canada's 700 MHz Auction set a limit of 500 bids. The limit was probably non-binding in the UK (with only 28 licences grouped into 4 categories, excluding low-power D licences), but it was likely binding in Canada (with 98 licences grouped into 56 categories).

Payment Rules: Allocation of the "Core Burden"

Core-selecting payment rules, in which winners are charged the opportunity costs of their winnings subject to the necessary "core" adjustments, have always been at the heart of the CCA design. In general, there are many payment rules that will satisfy this principle and it is important for the regulator to make an appropriate choice of core adjustment. While the core adjustment may appear to be an overly academic topic for regulators, it affects the fundamental fairness of the treatment of large versus small bidders.

To date, all implementations have used some variant of a "Nearest-Vickrey" payment rule that selects a unique set of payments from the minimum revenue frontier by minimizing the Euclidian distance to the Vickrey-Clark-Groves payments. Both Australia (2013) and Canada (2014 and 2015) made a simple improvement by employing a weighted version of the Nearest-Vickrey rule that minimizes a *weighted* Euclidian distance to the VCG payments. The strongest rationale for including weights is that, in auctions with regional licenses, the winning bidders might receive disparate winnings. Day and Cramton (2012) had motivated the Nearest-Vickrey rule, in part, by the fairness of bidders sharing the "core burden" equally. However, with disparate winnings, equal sharing might be unfair to those bidders with significantly smaller winnings.

The effects of the weighted Nearest-Vickrey rule can be illustrated using the actual data of Canada's 700 MHz Auction (see Table 3). In this auction, the VCG payments of two bidders, SaskTel and MTS, needed to be adjusted by a total of $10.376 mil ($CA) to satisfy the "blocking constraint" of their rivals. Under the Nearest-Vickrey rule, both bidders would have borne 50% of the core burden ($5.188 mil each). Instead, the weighted Nearest-Vickrey rule required MTS to bear approximately 54% of the burden, shifting the "core burden" towards MTS, which won greater value (as measured by the reserve price of the licences won).

Table 3: "Core Burden" in Canadian 700 MHz Auction (in mil $CA)

Bidder	**Base Vickrey Payment**	**Nearest-Vickrey Core Payment**		**Weighted Nearest-Vickrey Core Payment**	
		Split	Payment	Split	Payment
SaskTel	$2.755 mil	$5.188 mil	$7.943 mil	$4.802 mil	$7.557 mil
MTS	$3.198 mil	$5.188 mil	$8.386 mil	$5.574 mil	$8.772 mil

6 Conclusion

The CCA has quickly become one of the standard techniques in the toolbox of spectrum auction designers. Most academic effort is currently devoted to examining the major theoretical issues surrounding this auction format, while many of the questions important for practical applications are overlooked. Yet these small details can prove decisive for the overall success of the auction. Frequently, practitioners need to customize the basic CCA design in order to accommodate their specific objectives and unique environments. Fortunately, the CCA is a highly versatile design that allows integrating complex secondary objectives. Unfortunately, naïve integrations can introduce unintended consequences, damaging the auction's performance as a whole.

Recent CCAs have demonstrated that the details of implementing features even as standard as a set-aside can be far from trivial. And, when a regulator seeks to superimpose more complex side objectives on top of the standard CCA design, the nuances of the implementation can become critical. While the goal of accommodating additional features of the regulatory environment within a customised CCA is laudable, some of the custom design changes that have been implemented in various CCAs appear to contradict basic principles. For example, bidders in one auction were able to directly affect their own payments through a newly introduced option, contrary to the principles of opportunity cost pricing. Changes made to the winner determination process have created new gaming opportunities and modified activity rules have allowed some bidders to switch back and force between packages with very different degrees of competitiveness.

In this article, we have attempted to summarise and analyse some of the most common choices that a regulator needs to make when implementing the CCA. Regulators are encouraged to innovate on the basic CCA design to advance their novel objectives, but they are cautioned to pay heed of the basic principles underlying the CCA and thereby to avoid introducing novel design flaws.

University of Maryland, College Park

University of Colorado, Boulder

References

Ausubel, L.M. and Milgrom P. (2002). ‘Ascending auctions with package bidding’, *Frontiers of Theoretical Economics,* vol. 1(1), pp. 1-41.

Ausubel, L.M., Cramton P. and Milgrom P. (2006). ‘The clock-proxy auction: a practical combinatorial auction design’, in (P. Cramton, Y. Shoham, and R. Steinberg, eds.), *Combinatorial Auctions*, pp. 115-138, Cambridge: MIT Press.

Ausubel, L.M. and Baranov O. (2014). ‘Market design and the evolution of the combinatorial clock auction’, *American Economic Review: Papers & Proceedings,* vol. 104(5), pp. 446-451.

Ausubel, L.M. and Baranov O. (2016a). ‘Revealed preference and activity rules in auctions’, Working Paper, University of Maryland and University of Colorado.

Ausubel, L.M. and Baranov O. (2016b). ‘Vickrey-based pricing in iterative first-price auctions’, Working Paper, University of Maryland and University of Colorado.

Bichler, M., Shabalin P. and Wolf J. (2013). ‘Do core-selecting combinatorial clock auctions always lead to high efficiency? An experimental analysis of spectrum auction designs’, *Experimental Economics*, vol. 16(4), pp. 511–545.

Bichler, M., Goeree J., Mayer S. and Shabalin P. (2014). ‘Spectrum auction design: simple auctions for complex sales’, *Telecommunications Policy,* vol. 38(7), pp. 613–622.

Cramton, P. (2013). ‘Spectrum auction design’, *Review of Industrial Organization,* vol. 42(2), pp. 161-190.

Day, R.W. and Raghavan S. (2007). ‘Fair payments for efficient allocations in public sector combinatorial auctions’, *Management Science,* vol. 53, pp. 1389–1406.

Day, R.W. and Milgrom P. (2008). ‘Core-selecting package auctions’, *International Journal of Game Theory,* vol. 36, pp. 393–407.

Day, R.W. and Cramton P. (2012). ‘Quadratic core-selecting payment rules for combinatorial auctions’, *Operations Research*, vol. 60(3), pp. 588–603.

Janssen, M. and Karamychev V. (2013). ‘Spiteful bidding and gaming in combinatorial clock auctions’, Working paper, University of Vienna.

Levin, J. and Skrzypacz A. (2014). ‘Are dynamic Vickrey auctions practical?: properties of the combinatorial clock auction’, Working paper, Stanford University.

Nisan, N. (2006). ‘Bidding languages for combinatorial auctions’, in (P. Cramton, Y. Shoham, and R. Steinberg, eds.), *Combinatorial Auctions*, pp. 215-231, Cambridge: MIT Press.

Parkes, D.C. (2001). ‘Iterative combinatorial auctions: achieving economic and computational efficiency’, Doctoral Dissertation, University of Pennsylvania.