Combinatorial Auctioneering

Christoph Brunner, Jacob K. Goeree, Charles A. Holt, and John O. Ledyard*

July 26, 2006

Abstract

The simultaneous multi round (SMR) auction, introduced by the FCC in 1994, has been successfully applied in the sales of high-valued market licenses around the world. The FCC now contemplates setting a new standard that incorporates the possibility of package, or combinatorial, bids. This paper provides the first comprehensive laboratory test of several combinatorial auction formats that have been proposed in the recent literature. We find stark differences in terms of efficiencies and revenues, sometimes caused by seemingly minor design details. In general, however, the interest of policy makers in combinatorial auctions is justified by the laboratory data; there are simple package bidding formats that yield improved performance, especially in terms of seller revenue. We perform "stress tests" by considering environments where collusion is sustained in equilibrium in any of the (single-stage) multi-round formats. We introduce the two-stage *Anglo-Dutch combinatorial auction*, consisting of a multi-round clock stage followed by a single round of sealed bids, and find it effectively breaks collusion.

---

* Brunner, Goeree, and Ledyard: Division of the Humanities and Social Sciences, California Institute of Technology, Mail code 228-77, Pasadena, CA 91125, USA. Holt: Department of Economics, PO Box 400182, University of Virginia, Charlottesville, VA 22904-4182. We acknowledge partial financial support from the Federal Communications Commission (FCC contract 05000012), the Alfred P. Sloan Foundation, the Bankard Fund, the National Science Foundation (SBR 0094800), and the Dutch National Science Foundation (VICI 453.03.606). The conclusions and recommendations of this paper are those of the authors alone and should not be attributed to the funding organizations. We would like to thank Martha Stancill, Bill Sharkey, and Mark Bykowski for useful discussions. We are grateful to Raj Advani, Charlie Hornberger, Jou McDonnell, Anton Shuster, Walter Yuan, *Arzyx*, and *Eastek Design* for programming support, to Lauren Feiler, Joel Grus, Guido Maretto, Alan Miller, Laura Panattoni, Brian Rogers, and David Young for helping run the experiments, and to *Dash Optimization* for the free academic use of their Xpress-MP software.

I. Introduction

Simultaneous auctions for multiple items are often used when the values of the items are interrelated. An example of such a situation is the sale of spectrum rights by the Federal Communications Commission (FCC). If a telecommunications company is already operating in a certain area, the cost of operating in adjacent areas tends to be lower. In addition, consumers may value larger networks that reduce the cost and inconvenience of "roaming." As a consequence, the value of a collection of spectrum licenses for adjacent areas can be higher than the sum of the values for separate licenses. There can also be important synergies in the spectrum frequency dimension, where adjacent bands may improve capacity and reduce interference. For instance, in the FCC auction for air-to-ground communications frequencies in May 2006, a package of three bandwidth units sold for about 4.5 times as much as a single unit, and similar synergies were implied by unsuccessful bids. Value complementarities arise naturally in many other contexts, e.g. aircraft takeoff and landing slots, pollution emissions allowances for consecutive years, and coordinated advertising time slots. This paper reports a series of laboratory experiments to evaluate alternative methods of running multi-unit auctions, in both high and low-complementarities environments.

Various auction formats have been suggested for selling multiple items with interrelated values. The most widely discussed format is the simultaneous multiple round (SMR) auction, first used by the FCC in 1994. In the SMR auction, bidders are only allowed to bid on single licenses in a series of "rounds," and the auction stops when no new bids are submitted. To win a valuable package of licenses in this type of auction, bidders with value complementarities may have to bid more for some licenses than they are worth individually, which may result in losses when only a subset is won. Avoidance of this "exposure problem" may lead to conservative bidding, lower revenue, and inefficient allocations.

The obvious solution to the exposure problem is to allow bidding for packages of items. In such combinatorial auctions, bidders can make sure they either win the entire package or nothing at all. As a result, bids can reflect value complementarities, which should raise efficiency and seller revenue. Combinatorial bidding, however, may introduce new problems. Consider a situation in which a large bidder submits a package bid for several licenses. If other bidders are only interested in buying one of the licenses contained in that package, they might find it hard to coordinate their actions, even if the sum of their values is higher than the value of the package to the large bidder (the threshold problem). Thus, there is no clear presumption that package bidding will improve auction performance, and laboratory experiments may be used to evaluate alternative auction formats. Experiments can also be used to evaluate the extent to which auction procedures facilitate or deter bidder collusion. The next section summarizes the main features of the auction formats to be considered, and laboratory experiments to compare these formats are reported in the sections that follow.

II. Alternative Auction Formats

The various combinatorial auctions to be considered are best understood in terms of how they differ from the incumbent standard, the FCC's simultaneous multi-round auction procedure. Therefore, we will begin by explaining how the SMR auction was implemented in the experiments. Each auction consists of multiple rounds. Bidders have 40 seconds to submit their bids in each round. After time runs out or all bidders submitted their bids, the revenue maximizing allocation is determined. All bids from the current as well as previous rounds are taken into account. The bids that are part of the revenue maximizing allocation at the end of a round are called provisionally winning bids. Once no more bids are submitted, the auction stops and the provisionally winning bids become the final winning bids.

There are two constraints on bidding. The first constraint is an activity limit that determines the maximum number of different licenses for which a bidder can submit bids. Each bidder is assigned a pre-specified activity limit at the beginning of the auction. A bidder's activity limit falls if the number of submitted bids (plus the number of provisionally winning bids in the previous round) is less than the bidder's activity limit in the previous round. Activity is transferable, so a bidder with a limit of 3 could bid on licenses A, B, and C in one round and on licenses E, F, and G in the next round, for example. The second restriction is that each bid must exceed the previous high bid for that license by a specified bid increment. This requirement is a minimum, and new bids can exceed the "provisionally winning" bid by up to eight bid increments. The only exception to the increment rule is that the provisionally winning bidder is not required to raise that bid. Bidders can observe others' previous bids and can see which of those were provisionally winning. The effect of activity limits and bid increments is to force bids upward, although there are limited opportunities for withdrawing bids.[1] The auction stops after a round in which no new bids are submitted and no withdrawals occur.

This multi-round procedure can be adapted to allow for bids on both individual licenses and packages, and this approach has been shown to improve auction performance in some cases.[2] With package bidding, the relevant price of a license is not necessarily the highest bid on that license; indeed there may not even be a (non-package) bid on a particular license. One approach is to calculate the revenue-maximizing allocation of licenses after each round, and to use "shadow" prices that represent marginal valuations in terms of maximized revenue. Then the price of a package is the sum of the prices for individual items, and new bids in the subsequent round then have to improve on these prices by some minimum increment that depends on the size of the package. As with SMR, bidders are given the option of selecting one of a series of pre-

[1] As a partial remedy to the exposure problem, the FCC allows bidders to withdraw their provisionally winning bids in at most two rounds, at a penalty that equals the difference between their withdrawn bids and the subsequent sale price if that is lower. Porter (1999) reports laboratory data showing that the introduction of this withdrawal rule increases the efficiency of the final allocation as well as the seller's revenue.

[2] Rassenti, Smith and Bulfin (1982) first used experiments to compare the performance of sealed-bid auctions with and without package bidding. Ledyard, Porter, and Rangle (1997) provide data comparing several iterative processes. The combinatorial auction produces higher efficiencies in both designs.

specified higher increments. This approach, known as RAD (Resource Allocation Design) pricing, is due to Kwasnica, Porter, Ledyard, and DeMartini (2005). A variant of RAD pricing is being considered as an alternative to the SMR format currently used by the FCC.[3] One advantage of the RAD approach is that prices may convey information about how high a bidder must go to "get into the action" on a particular license or package.[4] The revenue maximization at the close of each round results in provisionally winning bids (on licenses or packages) and the associated RAD prices. The treatment of activity limits is analogous, with activity being calculated as the number of different licenses being bid for or being provisionally won in the previous round (separately or as part of a package). In order to prevent cycles, the bid increment is raised after a round in which revenue does not increase. As with SMR, the auction stops when no new bids are submitted, and the "provisional winning bids" for that round become the final winning bids (withdrawals are not needed with this format).

An alternative approach to the pricing problem is to have prices rise automatically and incrementally in response to excess demand, via a "clock" mechanism (Porter, Rassenti, Roopnarine, and Smith, 2003). In each round of the combinatorial clock (CC) auction, the price of a combination is the sum of the prices for each component, and bidders can indicate demands for individual items or combinations of items. If more than one bidder is bidding for an item, either separately or as part of a package, the clock price for that item rises by the bid increment. Otherwise, the price remains the same. There are no provisional winners, but other aspects of this auction are analogous, e.g., activity is defined in terms of the number of different licenses for which a bidder indicates a demand.[5] The auction typically stops when there is no longer any excess demand for any item.[6] One possible advantage of an incremental clock auction is that it prevents aggressive "jump bids," which have been observed by McCabe, Rassenti, and Smith (1988) in the laboratory and by McAfee and McMillan (1996) in an FCC auction. The clock-driven price increments may also alleviate the threshold problem of coordinating small bidders' responses to large package bids.[7]

---

[3] The variant considered by the FCC employs an "XOR" bidding rule, which means that each bidder can have at most one winning bid. Another difference concerns the pricing rule: in the FCC version, prices adjust slower in response to excess demand because they are "anchored" with respect to prices in the previous round. See Appendix D in the Goeree and Holt (2005) experiment design report for more details.

[4] Ideally, the license prices should represent the revenue value of relaxing the constraint that there is only one of each license. The discreteness in license definitions may, however, result in nonexistence of dual prices, and Kwasnica et al. (2005) propose a method of computing approximate prices.

[5] Note that Porter et al. (2003) did not use activity limits in their combinatorial clock auctions.

[6] When there is no more excess demand for any of the licenses but some are in excess supply, the revenue maximizing allocation is calculated using all bids in the current and previous rounds. If this process results in a failure to sell to the remaining bidder for an item, the clock is restarted to let that bidder have another chance to obtain the item.

[7] Ausubel, Cramton and Milgrom (2006) propose to add a second stage to the combinatorial clock auction. In the second stage, bidders would submit their valuations to a proxy agent who would bid up to that valuation in an ascending auction. The "clock-proxy" combinatorial auction has not yet been subjected to laboratory testing. Below we report results for a two-stage format that combines an ascending clock phase with a sealed-bid pay-your-bid phase, the "Anglo-Dutch" combinatorial auction.

Results of laboratory experiments suggest that these and other forms of package bidding may enhance performance measures, especially in environments with high complementarities.[8,9] In the Porter et al. (2003) experiment, for example, the combinatorial clock auction attained 100% efficiency in 23 sessions and 99% efficiency in two other sessions. Previous experiments have mainly focused on specific auction formats. This paper is an attempt to take the next step, i.e. to provide a systematic and parallel consideration of SMR and its most widely discussed alternatives. In addition, we consider a two-stage ("Anglo-Dutch") auction in which the ascending clock rounds are followed by a final round of sealed bids designed to break attempts at tacit collusion in an environment where each bidder has a natural "home market" as reflected by known value asymmetries.

III. A Preliminary Experiment: Threshold versus Exposure

It is instructive to begin with results from a simple three-license experiment to illustrate the exposure and threshold problems. There are three small bidders, each with an activity of 1 and an interest in a single license (A for bidder 1, B for bidder 2, and C for bidder 3). In addition, there is a large bidder 4 who is interested in all three licenses and has an activity limit of 3, see Table 1. The large bidder's values for individual licenses are randomly drawn from the interval from 4 to 7, with each integer value being equally likely. The large bidder's values for individual licenses increase as this bidder acquires more licenses according to a simple linear rule: when the large bidder acquires K licenses the value of each goes up by a factor $1 + \alpha\,(K\text{-}1)$, where we set the synergy factor $\alpha = 2$. In other words, individual values for this bidder are increased by a factor of 3 if two licenses are acquired, and the values are increased by a factor of 5 for a package of all three licenses. These high complementarities are intended to exaggerate the exposure problem for the large bidder in the absence of package bidding.

8 Banks, Ledyard and Porter (1989) proposed a different type of combinatorial auction, called Adaptive User Selection Mechanism (AUSM). In this auction, bidders can submit bids for individual licenses and packages in continuous time. A new bid becomes provisionally winning if revenue can be increased by an allocation that includes the new bid. Kwasnica, Ledyard, Porter and DeMartini (2005) compare RAD and AUSM to SMR in a laboratory setting. Efficiencies observed with RAD and AUSM are similar and higher than those for SMR, but revenue is higher in SMR since many bidders lose money due to the exposure problem. (If we assume that bidders default on bids on which they make losses and thus set the prices of such bids to zero, revenues are in fact higher under AUSM and RAD than under SMR.)

9 Charles River and Associates also developed a combinatorial auction, called Combinatorial Multi-Round Auction (CMA). In this auction, only bids that are sufficiently high allow bidders to maintain their activity. A bid is sufficiently high when it is at least 5% higher than the currently highest combination of bids that spans the same licenses. Banks, Olson, Porter, Rassenti and Smith (2003) ran an experiment to compare the CMA and SMR auction formats. They find that the CMA leads to more efficient allocations but less revenue since many bidders incur losses in their SMR auction experiments due to the exposure problem. Porter, Rassenti, Roopnarine and Smith (2003) also compare CMA to SMR auctions and also find that CMA tends to lead to more efficient allocations.

Table 1: Design for Preliminary Experiment

	License A	License B	License C	Activity	Synergy (α)
Bidder 1	[11,20]			1	0
Bidder 2		[41,60]		1	0
Bidder 3			[11,20]	1	0
Bidder 4	[4,7]	[4,7]	[4,7]	3	2

An asymmetry in small bidders' value distributions is used to create coordination problems when package bids are allowed. Small bidders 1 and 3 have values for their licenses of interest (A and C) that are drawn from the interval [11, 20]. In contrast, small bidder 2 has a value for B that is drawn from the interval [41, 60]. The value ranges were chosen so that, on average, the sum of the small bidders' separate values for A, B, and C would be approximately equal to five times the sum of the large bidder's values. In this manner, efficiency will sometimes require single license allocations to small bidders, and in other cases it will result in the allocation of two or all three licenses to the large bidder.

In the experiment, subjects were sorted into fixed groups of 4 bidders, with fixed roles for a series of 6 auctions, preceded by a practice auction. We used three different sequences of value draws, as indicated by the three rows of Table 2, both for RAD and SMR (without the possibility of bid withdrawals). Bidders' earnings are shown by group in Table 2, along with the theoretical earnings calculations for the optimal allocation under the assumption that losing bidders bid up to their license values.

Table 2: Actual and Theoretical Earnings (in Dollars)

	Small Bidders' Earnings			Large Bidder's Earnings		
Values	SMR	RAD	Theory	SMR	RAD	Theory
Sequence 1	243	3	16	-13	37	8
Sequence 2	72	45	37	17	20	12
Sequence 3	231	5	16	-20	57	19

Under SMR, the large bidder was unable to earn a profit in this harsh environment with high complementarities and strong competition from small bidders. Large bidders are frequently "exposed" when they compete aggressively for combinations of licenses and end up with a (low-value) subset. In turn, the risk associated with competing for packages caused some large bidders to drop out early, thus leaving substantial profits for small bidders who were able to close the auction without bidding close to their values (see the SMR column on the left in Table 2). Seller revenues were highly variable in this treatment, depending on how much risk the large bidder was willing to take. The distribution of revenues is shown by the upper cumulative revenue line in Figure 1: the horizontal axis corresponds to revenue as a percentage of maximum possible revenue (revenue when all bidders bid their values), whereas the vertical axis indicates the percentage of auctions with revenues below those levels. Note that in a few cases revenue is close to 100% (the large bidder is exposed and loses money) but in more than half of the cases revenue is less than 50% (the large bidder drops out early).

In a parallel series of auctions (based on identical value draws), we used RAD pricing and allowed for package bidding by the large bidder. It is apparent from Figure 1 that auction revenues were generally higher with package bidding, since the cumulative distribution of auction revenues (dark line) is lower, except at the far right. Recall that the theoretical earnings are calculated on the assumption that losing bidders allow bids to rise to the level of their values, which includes complementarities for the large bidders. This highly competitive bidding behavior might not be observed with package bidding if the small bidders are not able to coordinate bid increases that require bidder 2 to bid aggressively against the large bidder. The profits of the large bidder should be enhanced by any tendency for small bidders to drop out before prices approach their values, and it is apparent from the RAD column on the right side of the table that large bidders' earnings were higher than theoretical predictions for all 3 groups. Conversely, the effect of the threshold problem would be to reduce small bidders' earnings, which were generally, but not always, lower than predicted.[10]

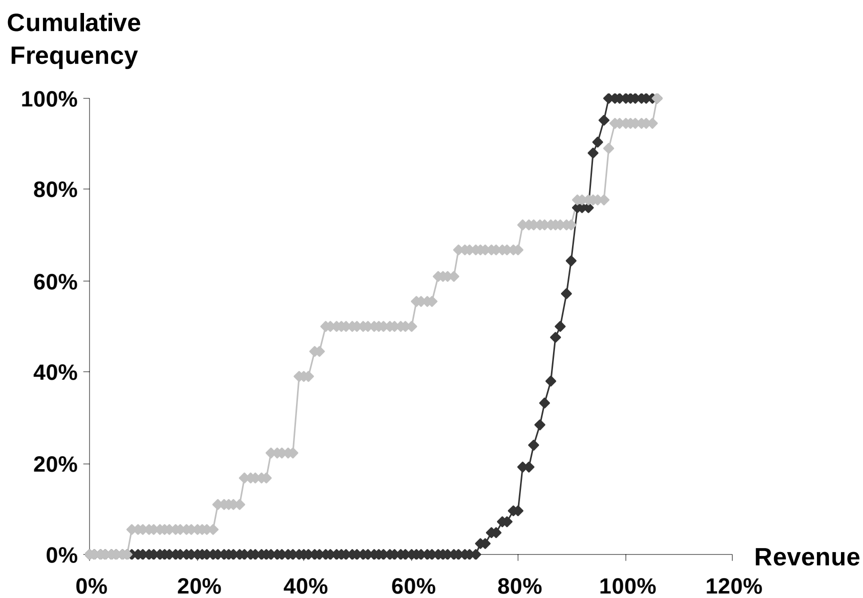

Figure 1: Cumulative Distribution of Auction Revenues
SMR (Upper Light Line) and RAD (Lower Dark Line)

[10] Another indication of the threshold problem in this environment can be found by looking at the extent to which small bidders lose activity and drop out before their bids reach their license values. With SMR, small bidders' bids reach 90 percent of license values for auctions where the large bidder wins, which roughly corresponds to bidding up to one's value minus one bid increment. In contrast, small bidders' bids stop at 80 percent of value, on average, for the RAD auctions where the large bidder wins.

A typical RAD outcome can be illustrated by the bid sequence from one of the groups, see Table 3. Columns 2-4 show the bids $b_1(A)$, $b_2(B)$, and $b_3(C)$ of small bidders 1, 2, and 3 for licenses A, B, and C respectively. Column 5 shows the large bidder's bid $b_4(ABC)$ for package ABC (the only package the large bidder bid on in this auction). The small bidders' values were 13, 50, and 11, but the large bidder provisionally won all three licenses in round 1 with a bid of 22 for the package ABC, as indicated by the last three columns of Table 3 that show the winners by round. The large bidder's package bid of 22 translated into RAD prices of 4, 10, and 8 shown in the price columns of the table. Second round bids by the small bidders were 6, 20, and 10 respectively, to which the large bidder reacted with a package bid of 58 in round 3. The resulting RAD prices of 19, 20, and 19 pushed bidders 1 and 3 out (i.e. their activities fell to 0) and only bidder 2 remained active. Over a series of rounds, this bidder raised the price on license B to 46 at which point the small bidders were provisionally winning at license prices of 6, 46, and 10. The large bidder then bid 68 for the package ABC, which resulted in RAD prices of 11, 46, and 11. Bidder 2 raised the price once more to 48, but the sum of the small bidders' bids (6 + 48 + 10 = 64) was not sufficient to unseat the package bid of 68. Notice that the symmetry of the RAD prices interacted with the asymmetry of small bidders' values to generate a threshold problem. Small bidders 1 and 3 were pushed out early in the auction and could no longer help bidder 2 in topping the winning package bid of 68, even though the sum of their values (13 + 50 + 11 = 74) exceeded this level. Bidder 2 is assigned license B in the optimal allocation, so in this example the inability for the small bidders to coordinate their bids lowers efficiency as well as revenue.

Table 3: Sequence of Bids in the RAD Auction

Round	$b_1(A)$	$b_2(B)$	$b_3(C)$	$b_4(ABC)$	p_A	p_B	p_C	W_A	W_B	W_C
1	2	10	8	22	4	10	8	4	4	4
2	6	20	10		6	20	10	1	2	3
3				58	19	20	19	4	4	4
4		28			15	28	15	4	4	4
5		32			13	32	13	4	4	4
6		38			10	38	10	4	4	4
7		46			6	46	10	1	2	3
8				68	11	46	11	4	4	4
9		48			10	48	10	4	4	4

The extreme complementarities and value asymmetries in this experiment illustrate the possible effects of the exposure problem without package bidding and the threshold problem with it. In the next section, we report a large-scale experiment with a more moderate valuation structure that was worked out in conversations with some FCC staff members as part of an earlier contract to compare SMR and a variant of RAD being considered for subsequent auctions.

IV. Main Experimental Design

Our main design involves groups of eight bidders and 12 licenses, a size that was selected to provide enough added complexity, while still permitting us to obtain sufficient independent observations for a broad range of auction formats and value structures. There are two types of bidders in this design: small "regional" bidders (labeled 1 through 6) and large "national" bidders (labeled 7 and 8). A graphical representation of bidders' interests is shown in Figure 2. Each diamond represents a different region, and the licenses along the center line (A, D, E, H, I, and L) are the ones of interest to the two national bidders. In the diamond shaped region on the far left, for example, the regional bidders, 1, 2, 5 and 6, are interested in licenses B and C, and in addition, each is interested in one of the licenses (A or D) that are targets for the two national bidders. Similarly, in the middle region, small bidders 1, 2, 3 and 4, are interested in licenses F and G, and each one is interested in one of the licenses (E and H) that are also of interest to the national bidders. The far-right diamond shaped region has a similar structure. Notice that each regional bidder has interests in two adjacent regions, e.g. the left and center diamonds for bidders 1 and 2.

Regional bidders can acquire at most three licenses, and complementarities occur only when licenses in the same region are acquired. For example, if bidder 1 wins the combination ABE, then the value synergies would only apply to A and B, which are in the same region in Figure 2. Since these value synergies do not apply across regions, the regional bidders face a more complex "fitting problem". National bidders can acquire up to six licenses and they have value complementarities for all six licenses in some treatments and for only four licenses in other treatments. The larger number of licenses subject to complementarities creates a larger exposure problem for the national bidders. The total number of possible allocations with this setup is 13,080,488.

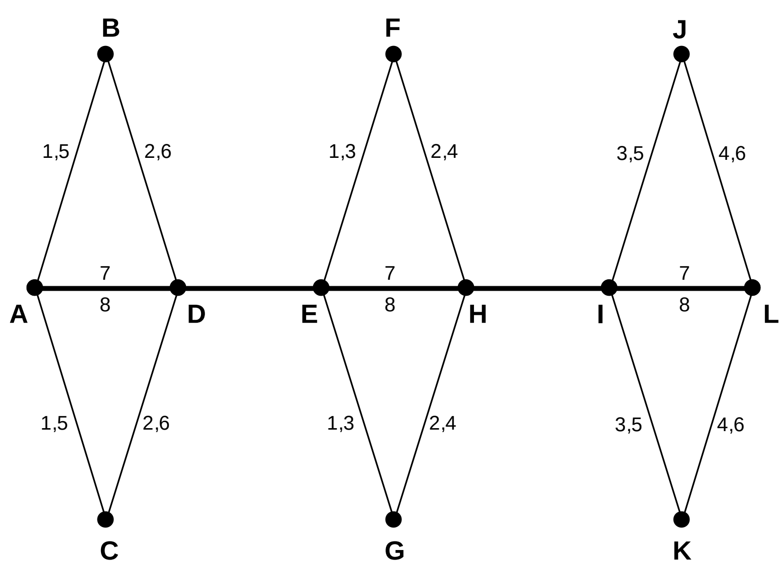

Figure 2: Eight-Bidder Design with Three Regions
Regional bidders (1-6) are interested in one side of one of two diamond-shaped regions.
National bidders (7-8) are interested in the middle line connecting all three diamond-shaped regions.

Auction Formats: We tested four auction formats that are described in detail in Appendices A-C. They include three combinatorial formats (FCC, RAD, and CC) and one non-combinatorial format (SMR). The CC and RAD formats were basically as described in the previous section, and the other two formats will be explained next.

The main modification of the basic SMR procedures described in the previous section is that bid withdrawals were permitted in at most two rounds of an auction. For example, a bidder who withdraws any number of bids in rounds 8 and 10 would not be able to make any withdrawals in subsequent rounds. If a withdrawn bid caused the final sale price to go down, the bidder had to pay the difference. If a license with a withdrawn bid went unsold, however, then the bidder was only responsible for 25% of the withdrawn bid, which represents a penalty intended to mimic the effect of having to pay the difference between a withdrawn bid and a lower sale price in a subsequent auction. A key feature of the withdrawal provisions is that the seller (FCC) becomes the provisionally winning bidder at the second highest bid (minus a bid increment), so that the person who originally made the second highest bid would be able to re-enter at that level if the bidder has activity and interest to do so. This provision can benefit a bidder whose interests have changed, perhaps to a different region.

The FCC format is a simple modification of the RAD auction described above, but with the requirement that bidders can have only one provisionally winning bid, an "XOR" rule. Notice that this rule facilitates the regional bidders' "choice of region" problem, since they can bid on packages from *both* regions knowing that at most one bid can be winning. In contrast, in the other formats, regional bidders have to change the contents of their bidding baskets if they want to switch to another region in response to price developments. Since XOR bidding typically calls for making bids on lots of combinations[11], the activity rule being considered by the FCC is based on the size of the largest package bid, so a bidder with activity 3 could bid on both ABD and ABC, for example, but not on ABCD. A final difference is that the FCC procedures use a modified method of approximating the dual prices associated with revenue maximization, and some "inertia" is built into the resulting price sequence.

For each auction format, the experiments cover four different treatments: high/low overlap in national bidders' interests (HO versus LO) and high/low complementarities (HC versus LC). For example, treatment HOHC has high overlap and high complementarities. We next describe the treatment variations in more detail.

Complementarities: Payoffs in the experiment were expressed in terms of "points," where each point was worth \$0.40 to subjects. The baseline draw distributions are uniform on the range [5, 45] for each license of interest to national bidders, and on the range [5, 75] for each license of interest to regional bidders. With high complementarities, the synergy factor (α) for national bidders was 0.2. Thus each license

[11] For example, the XOR rule means that a bidder who is interested in both licenses A and/or B must bid on A, B, and the package AB, since a bid on AB alone would preclude winning either license separately while bids on A and B only would preclude winning the package.

acquired by a national bidder goes up in value by 20% (with two licenses), by 40% (with three licenses), by 60% (with four licenses), by 80% (with five licenses) and by 100% (with all six licenses). With low complementarities, these numbers are 1%, 2%, 3%, 4% and 5%, corresponding to $\alpha = 0.01$. With high complementarities (HC), each license acquired in the same region by a regional bidder goes up in value by 12.5% (with two licenses in the same region), and by 25% (with three licenses in the same region), so $\alpha = 0.125$. With low complementarities, these numbers are 1% and 2% for regional bidders. These minimal complementarities in the LC treatment allowed us to maintain parallelism in instructions and procedures.

Overlap: With high overlap (HO), each national bidder, 7 and 8, has value draws from the same distribution for all six licenses on the base of Figure 2, and the complementarities apply equally to all six licenses. In this sense, each national bidder is equally strong across the line. With low overlap (LO), national bidder 7 only receives complementarities for the four licenses on the left side of the base (A, D, E, and H). Conversely, national bidder 8 receives complementarities for the four licenses on the right side (E, H, I, and L). Thus with high complementarities and low overlap, each national bidder has a natural focus of interest that only partially overlaps with the other national bidder's area. One issue of interest is whether this type of partial separation may yield tacit collusion and less aggressive bidding in the center.

Treatment Structure: The two-by-two treatment design yields four treatments for each of the four auction formats, for a total of 16 treatments. We used the same value draws across auction formats and treatments so that differences cannot be attributed to specific sequences of value draws. Each session consisted of one or two practice auctions and a series of six auctions for cash payments. The treatment and auction type was unchanged for all auctions in a session, but the randomly generated value draws changed from one auction to the next. In addition, we used new sequences of random draws for each of three "waves" of 16 sessions that spanned all treatments.

Subjects and Sessions: Before conducting the sessions that form waves 1-3, we trained over 128 Caltech subjects in 16 sessions of eight people. These inexperienced sessions ("wave 0") involved both SMR and combinatorial auctions and were conducted to familiarize subjects with the auction software and bidding environment.[12] For these inexperienced sessions, we promised to pay each person a $60 bonus (in addition to other earnings) if they returned three more times. This decision to use experienced bidders was based on the complexity of the auction formats and on earlier pilot experiments. For the subsequent data analysis, only the data from waves 1-3 but not from wave 0 is used. In waves 1-3, earnings averaged $50 per person per session, including $10 show-up fees and bonuses, for sessions that lasted from one and a half to two hours, depending on the

[12] The experiments were run using *jAuctions*, which has been developed at Caltech by Jacob Goeree. The *jAuctions* software consists of a flexible suite of Java-based auction programs designed to handle a wide range of auction formats and bidding environments, including combinatorial auctions with bid-driven or clock-driven prices, private and common valuations, etc. Instructions, which are available on request, were structured around relevant screen shots of the *jAuctions* program.

number of auctions. In total, there were 16 training sessions and 48 sessions (3x16) with experienced subjects, each involving a group of eight subjects.

V. Results: Comparing Alternative Auction Formats

One way to measure market efficiency is to divide the sum of all bidder values for licenses they won, the actual surplus (S_{actual}), by the maximum possible surplus ($S_{optimal}$). It is well known that this simple efficiency measure may be difficult to interpret. For example, adding a constant to all value amounts will tend to raise this efficiency ratio, since efficiency losses are affected by differences in valuations, not absolute levels. A more natural measure of efficiency is calculated on the basis of the difference between the actual surplus and the surplus resulting from a random allocation (S_{random}), this being normalized by the maximum such difference.

$$efficiency = \frac{S_{actual} - S_{random}}{S_{optimal} - S_{random}} * 100\%$$

The value of a random allocation can be computed by taking the average of the surplus over all possible allocations, of which there are 13,080,488 in total for the design in Figure 2.[13] This definition of efficiency measures how much the auction raises surplus relative to a random allocation mechanism. In the analysis that follows, we will use these normalized efficiency measures.

Similarly, revenues will be measured as the difference between actual auction revenue and the revenue from a random allocation in which bidders pay their full values for all licenses and packages they receive ($R_{random} = S_{random}$). This difference is then divided by the difference between the maximum possible revenue and the revenue from a random allocation.

$$revenue = \frac{R_{actual} - R_{random}}{R_{optimal} - R_{random}} * 100\%$$

Since $R_{random} = S_{random}$ and $R_{optimal} = S_{optimal}$, the denominators of the normalized efficiency and revenue measures are equal, and the normalized sum of bidders' profits is simply equal to the difference between efficiency and revenue:

$$profit = \frac{S_{actual} - R_{actual}}{S_{optimal} - R_{random}} * 100\% = \frac{\sum_i \pi^i_{actual}}{S_{optimal} - S_{random}} * 100\%$$

[13] Without the restriction that regional bidders can acquire at most three licenses, the total number of allocations would simply be $4^{12} = 16,777,216$.

All efficiency, revenue, and profit measures reported below are normalized in this manner for the specific value sequences used in each auction for each of the three waves of sessions with experienced bidders.

Efficiency: Package bidding is designed to help bidders avoid the "exposure problem" of bidding high for licenses with high complementarities. As expected, switching from SMR to a combinatorial format raises efficiency in the high complementarities treatments. Combining low and high overlap treatments and data from all three waves, the average efficiency in SMR is 84% while it is 90%, 90% and 91% in FCC, CC, and RAD respectively. These large differences between SMR and the combinatorial formats occur for both of the high complementarities treatments, as can be seen from the left side of Figure 3. In this and subsequent figures the color-coding is as follows: from light to dark the bars correspond to SMR, CC, FCC, and RAD respectively.

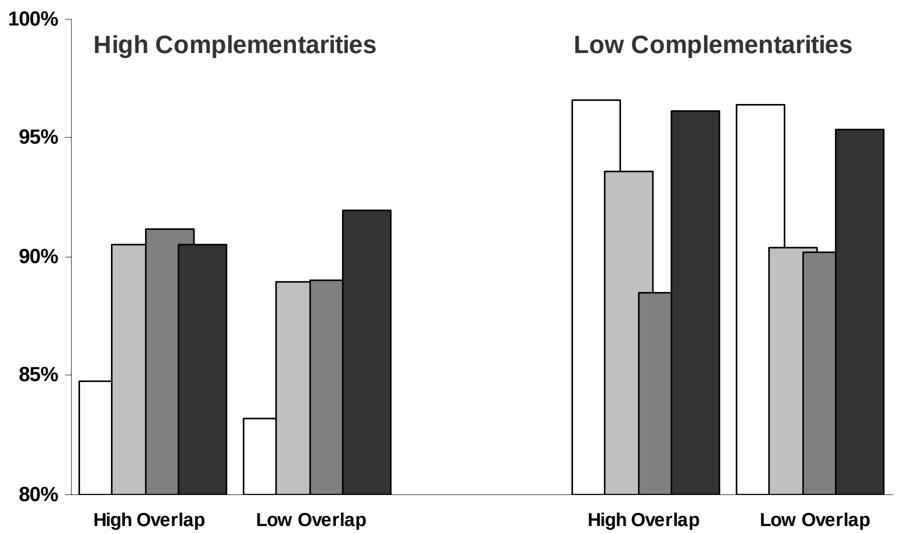

Figure 3: Efficiency by Auction Format
The bars from light to dark (left to right) correspond to SMR, CC, FCC and RAD respectively.

In contrast, the switch to combinatorial auctions reduces efficiency when complementarities are minimal (our "low complementarities" treatment). The efficiency levels are now 97% for SMR and 89%, 92% and 96% for FCC, CC, and RAD respectively. Again, this difference shows up in both LC treatments shown on the right side of Figure 3. Result 1 summarizes our findings, where we use the following notation: $\sim$ implies a pair wise difference is not significant, $\succ^{*}$ indicates significance at the 10% level, $\succ^{**}$ indicates significance at the 5% level, and $\succ^{***}$ indicates significance at the 1% level.

Result 1: *With high complementarities, efficiency levels are highest for the three combinatorial formats and are ranked*

$$RAD \sim FCC \sim CC \succ^{**} SMR$$

With low complementarities, efficiency levels are ranked

$$RAD \sim SMR \succ^{**} CC \sim FCC$$

Pooling the low and high complementarities treatments, efficiency levels are ranked

$$RAD \succ^{*} CC \sim FCC \sim SMR$$

Support: Session averages are grouped by wave and auction format in Appendix D. For example, consider the efficiencies for the HC treatments (pooling high and low overlap) shown in the eight columns on the left side of Appendix D (top three rows). It is important to compare the auction formats for the same wave, since the valuation draws change from one wave to another. All six of the paired comparisons for the HC treatments show higher efficiencies for any of the package bidding auctions other than SMR. This effect is significant using a Wilcoxon matched-pairs signed-rank test (p = 0.03). These results are generally reversed with low complementarities, where all paired comparisons between SMR and CC and FCC go in the opposite direction (higher efficiency for SMR): this effect is significant (p = 0.03). The only combinatorial auction that is not statistically different from SMR in terms of efficiency is RAD (p=0.41). When pooling the data from the low and high complementarities treatments, RAD is more efficient than FCC (p = 0.02), CC (p = 0.09), and SMR (p = 0.09). There are no significant differences between SMR, FCC, and CC. █

One reason why SMR leads to low efficiencies with high complementarities is the incidence of unsold licenses, which happens with rates of 4% and 7% in the high and low overlap treatments respectively, see Appendix D. Unsold licenses typically result from withdrawals late in the auction when a bidder realizes that it will not be possible to obtain the value synergies associated with multiple licenses. After a withdrawal, recall that the seller becomes the provisional winner at the second highest bid, and the person who made that bid previously may have lost activity or interest in that license, which causes it to go unsold. Withdrawals are not permitted in the combinatorial auctions, where the exposure problem is addressed directly by allowing package bids, so these auctions do not result in unsold licenses. The difference between SMR and any of the combinatorial formats in terms of license sales rates is significant with a Wilcoxon matched-pairs signed-ranks test (p = 0.05).

Revenues: Figure 4 shows the revenues by auction format and treatment averaged across sessions (session averages for each parameter/experience wave can be found in rows 4-6 of the table in Appendix D). What is obvious from Figure 4 is that the combinatorial clock auction extracts more rents for the seller in all treatments, even when it is less efficient than other formats.

Result 2: *Revenues are highest for the combinatorial clock auction and are ranked*

$$CC \succ^{***} RAD \sim FCC \sim SMR$$

Support: There are three rows in the Revenue section of Appendix D, one for each wave of parameter values. In each row, there are 4 paired comparisons between CC and a particular alternative format, so overall there are 12 paired comparisons. The CC provides higher revenue in all 12 pair-wise comparisons with each of the alternatives, except for RAD where CC yields higher revenues in 11 of 12 cases. These comparisons are significant using a Wilcoxon matched-pairs signed-rank test ($p = 0.001$). Basically, CC is higher than the others with both low and high complementarities ($p = 0.03$), except for the comparison with RAD at LC ($p = 0.06$). Revenue under RAD is border-line significantly higher than FCC when we pool all data ($p = 0.109$). Revenue under RAD is not significantly different from SMR, and SMR and FCC raise the same revenues. ■

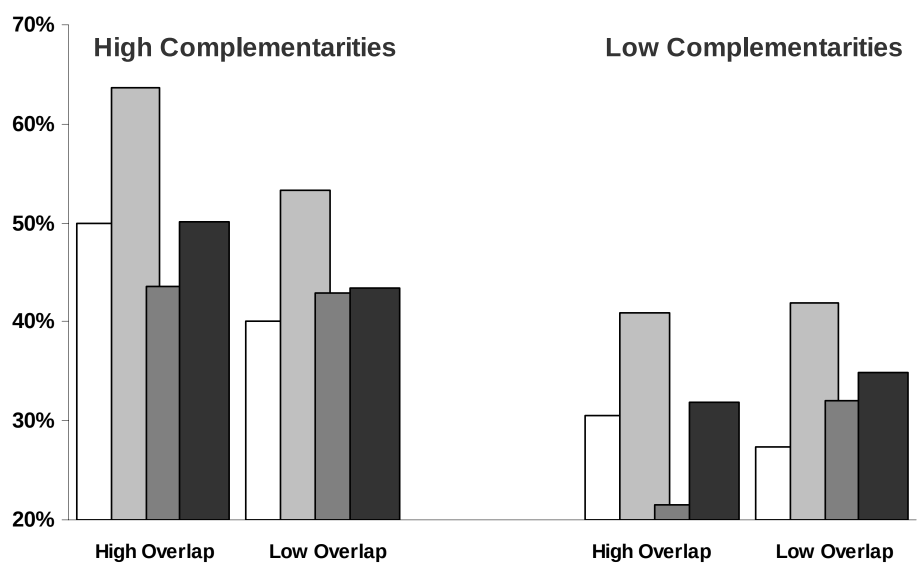

Figure 4: Revenue by Auction Format
The bars from light to dark (left to right) correspond to SMR, CC, FCC and RAD respectively.

Profits: Figure 5 shows bidders' profits by auction format and by treatment. The ability to bid for combinations allows bidders to bid high on packages and avoid the exposure problem, an effect that is mainly relevant with high complementarities. But if all bidders bid higher, the effect on bidder profits is unclear.

Result 3: *Bidders' profits are lowest in the combinatorial clock auction and are ranked*

$$RAD \sim FCC \sim SMR \succ^{***} CC$$

Support: Normalized profits are calculated as the differences between entries in the efficiency and revenue rows of Appendix D. With three waves and four treatments, there are 12 paired profit comparisons between CC and a particular alternative format, and the CC provides lower profits in all 12 pair-wise comparisons with SMR, and for 11 of the 12 comparisons with RAD and FCC. These comparisons are significant using a Wilcoxon matched-pairs signed-rank test (p = 0.001). Averaged over treatments, profits for CC are 40% and profits for the other formats are all in a narrow range from 53-55%. ■

The exposure problem can be alleviated to some extent by the (limited) bid withdrawal provisions built into the SMR bidding rules under consideration. In this manner a bidder may compete aggressively for a package and then decide to withdraw, paying a penalty equal to the difference between the withdrawn bid and the final sale price if it is higher. Withdrawals are more frequent (and the associated penalties higher) with high complementarities, as would be expected. The effect of withdrawal penalties on bidders' profits was generally small.[14]

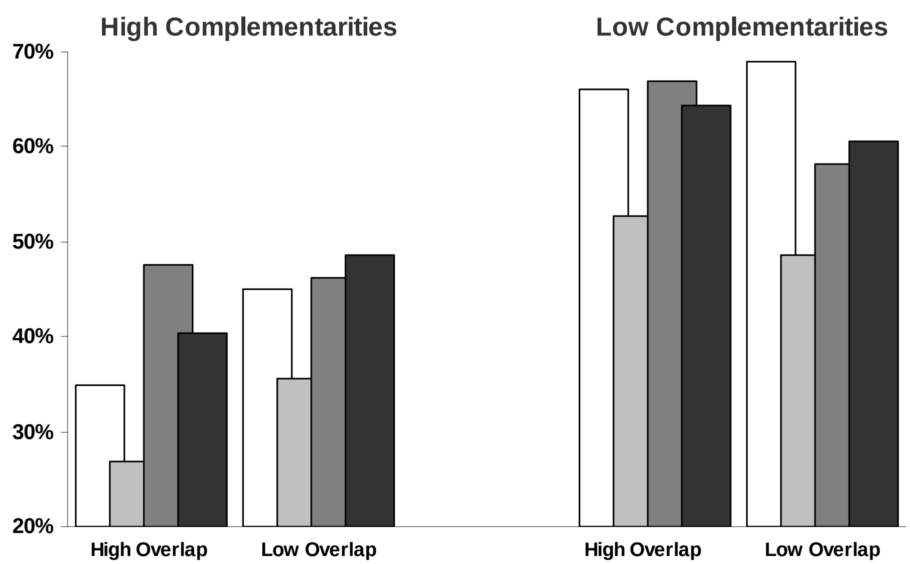

Figure 5: Bidders' Profits by Auction Format
The bars from light to dark (left to right) correspond to SMR, CC, FCC and RAD respectively.

[14] National bidders' penalties averaged 2% in the HC treatments and were negligible in the LC treatments. Regional bidders' penalties averaged 1% in the HC treatments and were negligible in the LC treatments.

Summary: Pooling data across treatments and sessions, the revenue and efficiency results by auction format are given in Table 4. In terms of seller revenues and buyer profits, the combinatorial clock auction is best for the seller and worst for the bidders, but these results are not caused by bidder losses, which are not present in the CC sessions (see the losses rows for Nationals and Regionals in Appendix D). In a comparison with the other formats, the FCC auction with XOR bidding is the worst from the seller's point of view (lowest revenue and efficiency), and it is the best from the bidders' point of view (profits). Our conjecture for why the FCC auction performs the worst in terms of efficiency is that, in the presence of minimal complementarities, the requirement that bidders can only have one bid accepted (the XOR rule) may reduce efficiency, since bidders have to bid on many combinations of licenses to find all possible efficiency gains. (Even though our design, in which regional bidders face a "choice of region" problem, favors the XOR rule.) RAD, in contrast, more or less reduces to SMR with low complementarities while it enables bidders to extract the extra efficiency gains when complementarities are high. Another consideration may be that the inertia in the FCC price adjustment algorithm could exaggerate the threshold problem, since attempts to unseat large package bids may have delayed effects due to inertia. This conjecture might be evaluated by running the asymmetric design discussed in Section III for the FCC format.

Table 4: Summary Statistics by Auction Format

	SMR	CC	RAD	FCC
Efficiency	90.2%	90.8%	93.4%	89.7%
Revenue	37.1%	50.2%	40.2%	35.1%
Profits	53.1%	40.6%	53.2%	54.6%

V. Extensions: Combating Collusion

Although many of our experienced subjects repeatedly participated in similar types of auctions, collusion did not play a role with our main design.[15] One reason could be that the market is relatively "thick," i.e. all licenses in Figure 2 are of interest to four bidders. Here we consider a "thin" market design where three bidders compete for three licenses, A, B, and C. All three bidders are interested in all licenses, but each bidder has a preferred license:[16] bidder 1 has a high value for A and low values for B and C, bidder 2 has a high value for B and low values for A and C, and, finally, bidder 3 has a high value for C and low values for A and B, see Table 5. High values are 19, 20, or 21 (with equal probabilities) and low values are 1, 2, 3, 4, or 5 (with equal probabilities). In addition, bidders have synergistic values for packages. If a bidder acquires two licenses then the

[15] Even though on more than one occasion subjects publicly announced they would collude just before the experiment started.
[16] See Li and Plott (2005) for a related setup and the importance of creating "preferred licenses" for the occurrence of collusive behavior.

value of each license goes up by 50% and if a bidder acquires all three licenses then the value of each license goes up by 100%. For example, if bidder 1 has a value of 20 for A and values of 3 and 4 for B and C respectively, then the value for the package AC is 36 and the value for ABC is 54.

Table 5: Design for the Collusion Experiment

	License A	License B	License C	Activity	Synergy (α)
Bidder 1	[19,21]	[1,5]	[1,5]	3	0.5
Bidder 2	[1,5]	[19,21]	[1,5]	3	0.5
Bidder 3	[1,5]	[1,5]	[19,21]	3	0.5

To avoid repeated game issues (i.e. collusion across auctions) we used a "round robin" design. For each session, we invited nine people to the lab who competed in groups of three. The matching was done such that no two bidders were in the same group more than once (and subjects were informed of this feature). As a result, we could only do four auctions with a group of nine subjects.

In any of the multi-round auctions discussed above, there are two natural outcomes with this setup: a competitive outcome where bidders try to acquire the largest and most valuable ABC package and a collusive outcome where bidders only bid for their preferred licenses. Both can be sustained in equilibrium.[17]

Proposition 1: *For the design in Table 5, multi-round auctions support both competitive and collusive outcomes.*
(i) In the collusive equilibrium, bidders bid only on their preferred license and the auction ends immediately (at initial prices). If a bidder deviates, others rotate their bids among the licenses such that all three license prices rise at the same pace.
(ii) In the competitive equilibrium, all three bidders compete for the ABC package and are willing to bid up to their package values.

Proof: In the collusive equilibrium, each bidder earns the value for the preferred licenses minus the initial price, which is at least 19 – 1 = 18. A deviation will trigger others to drive up the prices for all licenses at the same rate. In equilibrium, a bidder will bid at least 19 for the preferred license so others can, without risk, bid any license price up to 19. The auction ends when either the deviant bidder "folds" by bidding only for their preferred license (in which case they would have been better off doing so at the start of the auction) or the deviant bidder keeps bidding on the ABC package until the single-license bidders are "pushed out" of the auction. Since the maximum value of the ABC package is 62 (when the deviant bidder's license values are 21, 5, and 5 respectively) this scenario results in a profit of at most 62 – 57 = 5, which is again less than would have

[17] Other equilibria may exist but here we focus on the two equilibria that are most salient in our data.

been obtained by bidding on a single license. Hence, the collusive strategy to bid only for the preferred license (and drive up prices at an equal rate when a deviation occurs) constitutes an equilibrium. Finally, when others are willing to bid up to their values on the ABC package, it is clearly in a bidder's interest to bid on this package since the highest possible value of a subset (1.5 * 26 = 39) is less than the lowest possible value of the ABC package (2 * 21 = 42). ∎

Note that bidders' profits will be low and revenue will be high in the competitive equilibrium. Even when a bidder has the highest possible values, their value for the ABC package of 62 is only 10 higher than its expected value of 52. In contrast, when bidders compete only for their preferred licenses the auction closes immediately at low prices and bidders' profits will be close to 20.

We ran three sessions using the combinatorial clock auction, as this format raised the highest revenues with our main design, with one important change: bidders are no longer informed about others' bids. Each session results in 12 observations (three groups and four periods) for a total of 36 observations. The cumulative distribution of revenues for the combinatorial clock auction is given by the top line in Figure 6. Revenue levels (measured as a percentage of the value of the best possible allocation) are shown on the horizontal axis while the percentage of auctions that yielded a certain revenue level or less is shown on the vertical axis.

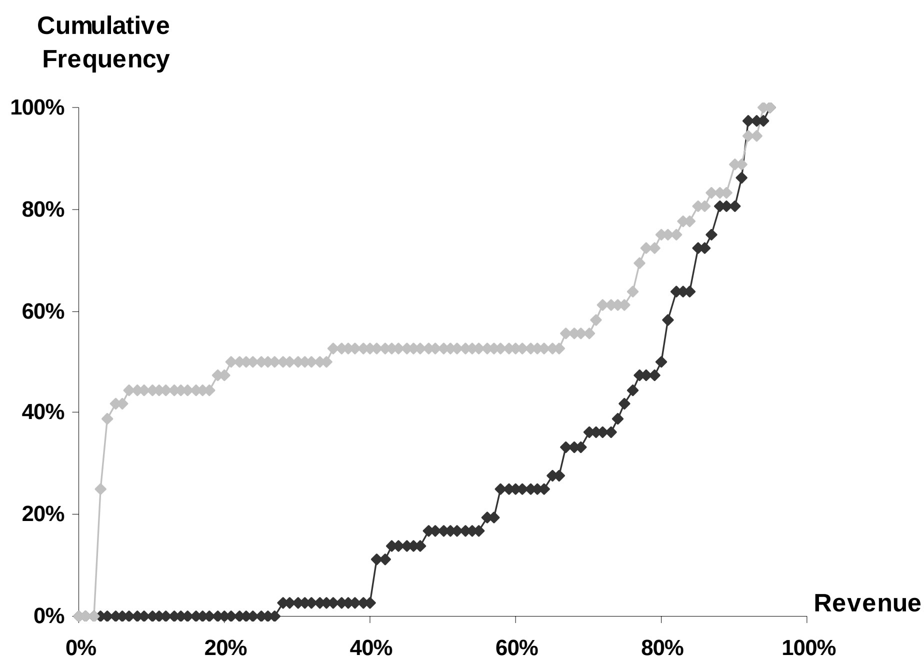

Figure 6: Cumulative Distribution of Auction Revenues
Combinatorial Clock (Upper Light Line) and Anglo-Dutch (Lower Dark Line).

From the figure it is apparent that both the competitive and collusive outcomes occur in the combinatorial clock auction. In more than 40% of the cases, revenue is less than 10%, i.e. the auction closes immediately. In more than 50% of the cases, revenues are less than 40%. In the remaining cases, revenues are 60% and above, in line with the competitive outcome. To summarize, the distribution of revenues is bi-modal and both the competitive and collusive outcomes are observed. The average revenue is low (44.3%) and its standard deviation is high (39.1%), see Table 6.

One possible way to break collusion is to append a second ("sudden death") stage to the combinatorial clock phase. Several implementations are possible, e.g. the second stage could be a proxy auction that has a second-price flavor (Ausubel, Cramton, and Milgrom, 2004). Here we consider a first-price or "pay-your-bid" stage that follows a simplified combinatorial clock stage: the Anglo-Dutch combinatorial auction (Klemperer, 2002).[18] As in the clock auction, all three bidders start with an activity level of 3 in the first stage and prices on items for which there is excess demand rise. Bidders do not observe others' bids and hence do not know their activity levels. But once the group activity falls below 6, the second stage is initiated in which bidders can make a sealed bid on as many licenses (separately or in a package) as their current activity allows for (in other words, bidders' activity limits are carried over from the first to the second stage).

There are no restrictions on the sizes of the bids. In other words, even when license prices have risen to, say, 15 or more in the first stage, bidders may place bids as low as 1 in the second stage. However, bidders know that all bids from both stages of the auction are used to determine the final winners, so a low second-stage bid is likely to become a losing bid. Allowing for flexibility in the second-stage bids alleviates the problem that small bidders are "pushed out" of the auction by a large first-stage package bid.

Proposition 2: *In the Anglo-Dutch combinatorial auction, the collusive outcome of Proposition 1 cannot be sustained in equilibrium.*

Proof: The collusive outcome requires that all three bidders only bid for their preferred license in the first stage, which would stop immediately as a result. In the second stage they would bid 1 for their preferred licenses. However, if bidders 2 and 3 bid this way, bidder 1 has an incentive to bid on the ABC package in stage 1, which still ends immediately (since the group activity level of 5 is below the threshold level of 6). In the second stage, bidder 1 then wins the high-value ABC package at a price of 3. █

Table 6: Comparison of the Anglo-Dutch and Combinatorial Clock Auction

	Combinatorial Clock	Anglo Dutch
Efficiency	98.2% (4.7%)	94.7% (6.4%)
Revenue	44.3% (39.1%)	74.7% (18.2%)

[18] Klemperer (2002) first proposed the Anglo-Dutch format in the context of single-license auctions.

We ran three sessions with groups of 9 bidders using the Anglo-Dutch combinatorial format. The distribution of revenues for the Anglo-Dutch auction is shown as the bottom line in Figure 6. With this auction format, less than 20% of the outcomes result in revenues of 50% or less. Indeed, more than 50% of the auctions generate revenues of 80% or more. The average revenue raised by this format is 74.7% with a standard deviation of 18.2%, see Table 6. To summarize, the Anglo-Dutch combinatorial auction is effective at breaking collusion and results in higher, less variable revenues.

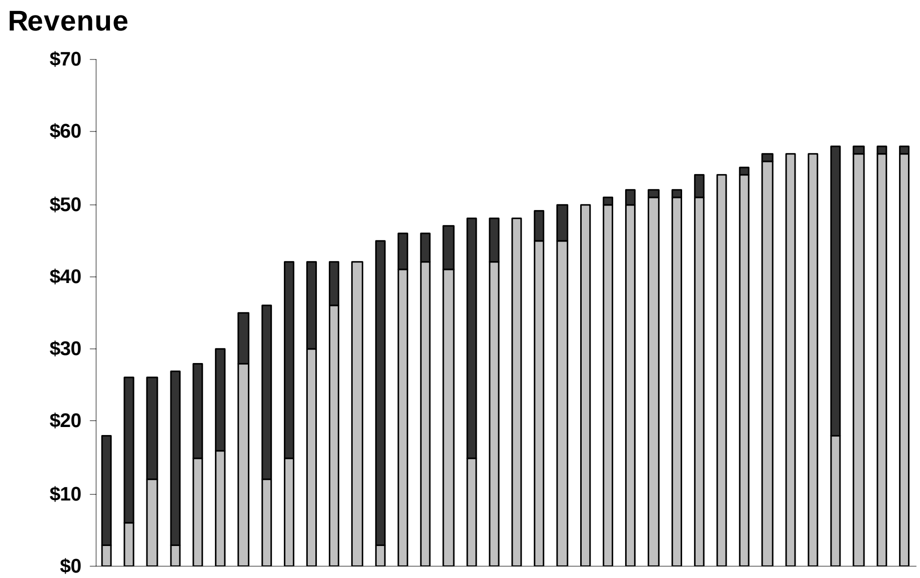

Figure 7: Revenues in the Anglo-Dutch Combinatorial Auction
Revenue after First Stage (Lower Light Bar) and Revenue after Second Stage (Upper Dark Bar).

It is interesting to consider in detail some of the bid sequences to understand why the Anglo-Dutch performs so much better than the clock auction in this environment. Figure 7 displays the revenues (in dollars) for all 36 auctions conducted with the Anglo-Dutch format. The light bottom part of a revenue bar represents the amount the three licenses would sell for if the auction closed without a second stage, i.e. the sum of the three clock prices at the end of the first stage. The dark top part of a revenue bar shows how much is added in the final "shootout" stage. There are two kinds of "almost collusive outcomes" in our data. In one example, two out of three bidders started out by bidding on only their preferred item so the first stage ended with clock prices equal to 1 for all three licenses. However, one of the bidders had an activity limit of 3, and this bidder put in a package bid on ABC of 45 to ensure it would win. In another example, bidding in the first stage resulted in prices of 6 on all three licenses. At this point all three bidders had activity limits of 1 but none could be sure about others' activity limits (except that the sum of the other two bidders' activity limits was 4 or less). The bids in the second stage were 17, 12, and 6 for a total profit of 35 (twice is much as it was after the first stage). Most bid sequences, however, show that bidders were unwilling to give up activity because they realized this would adversely affect their chances in the second stage. In 22 out of the 36 auctions we ran with this format, license prices at the end of first stage were somewhere in the 14-20 dollar range. To summarize, the Anglo-Dutch auction generally creates sufficient strategic uncertainty for bidders to compete in the first stage, driving clock prices up to near competitive levels. In situations of almost collusion in the first stage, bidders' uncertainty about others' bid possibilities generally adds substantial amounts to the seller's final revenues.

VI. Conclusions

The simultaneous multi-round auction is considered to be a remarkably successful application of game theory, with careful attention to the details of implementation by policy makers. This auction format is currently used around the world, and government officials in other agencies now routinely consult the FCC on auction design matters. Concerns about bidder collusion and the effects of value complementarities have convinced many people that new procedures need to be developed and tested. In particular, FCC officials recently arranged for a systematic experimental study of the effects of allowing package bidding with feedback that includes "shadow price" information about license values ("RAD prices") as they evolve in a multi-round auction. That study (Goeree, Holt, and Ledyard, 2006) involved a large-scale comparison of the incumbent SMR format with a variant of a RAD auction that was designed by other consultants, which we call the FCC format. This paper extends that study by including two other formats, the combinatorial clock (CC) and the original RAD auction that does not restrict bidders to one winning bid, i.e. without the "XOR" bid rule. Clock auctions are being widely discussed, in part because of near perfect performance in some prior laboratory tests (Porter et al., 2003). In addition, we introduce a two-stage Anglo-Dutch combinatorial auction that is designed to break collusion in environments where known value asymmetries create natural market divisions.

The experiments were conducted with a common *jAuctions* bidder interface and parallel sets of value draws, for an array of structural and auction format treatments. The combinatorial auction procedures used (RAD, FCC, and CC) all result in higher efficiency than the currently used SMR procedure when value complementarities are present. It is important to emphasize that value complementarities are not just a theoretical possibility; a package of three bandwidth segments sold for about five times as much as a single segment in a recent FCC auction that offered a very limited menu of pre-specified package bidding options. Complementarities are almost surely significant for other potential applications of package bidding such as emissions permits for successive years. The efficiency enhancing effects of combinatorial bidding, however, was mitigated or reversed in low-complementarities treatments for all combinatorial auction formats, with the exception of RAD.

The most dramatic auction failures in other countries have involved low revenues, and revenue generation seems to be the chief concern of policy officials. In this dimension, the combinatorial clock clearly outperforms the SMR and other combinatorial formats, in all treatments. The revenue gain for the CC is dramatic given the moderate nature of the experiment design, and the potential gains from implementing this procedure in high-stakes spectrum auctions are likely to be quite large, in our opinion. Low-revenue outcomes in spectrum auctions abroad have also been attributed to tacit collusion when some firms have an incumbency advantage, and our experiments provide evidence that this collusion can be negated by adding a "final shootout" stage to the auction. In the Anglo-Dutch combinatorial auction, the ascending bid stage (implemented as a clock auction) is terminated when activity falls below a pre-specified threshold, and then bidders submit final bids on licenses or packages, with the winners being determined by a revenue-maximization procedure. This procedure yields sharp increases in revenues in a setting where tacit collusion in a multi-round ascending-bid auction is a common outcome.

Taken together, these results suggest a flexible approach to auction design, with combinatorial procedures being used where value complementarities are high and a two-stage approach in "thin" auctions where tacit collusion is a consideration. Of the combinatorial procedures, RAD provides a nice balance in terms of efficiency and revenue while the combinatorial clock seems to be the best choice when revenue generation is the primary consideration, which is likely the case in most applications.

References

Ausubel, L. M., P. Cramton, and P. Milgrom (2006) "The Clock-Proxy Auction: A Practical Combinatorial Auction Design," in *Combinatorial Auctions*, eds. P. Cramton, Y. Shoam, and R. Steinberg, MIT Press, p. 115-138.

Banks, J. S., Ledyard, J. O. and Porter, D. P. (1989) "Allocating Uncertain and Unresponsive Resources: An Experimental Approach," *Rand Journal of Economics*, 20(1), 1-25.

Banks, J., Olson, M., Porter, D., Rassenti S. and Smith V. (2003) "Theory, Experiment and the Federal Communications Commission Spectrum Auctions," *Journal of Economic Behavior and Organization*, 51, 303-350.

Goeree, J. K. and T. Offerman (2004) "The Amsterdam Auction," *Econometrica*, 72(1), January, 281-294.

Goeree, J. K., T. Offerman, and R. Sloof (2005) "Demand Reduction and Preemptive Bidding in Multi-Unit License Auctions," Caltech Working paper.

Goeree, J. K. and C. A. Holt (2005) "Comparing the FCC's Combinatorial and Non-Combinatorial Simultaneous Multi-Round Auctions: Experimental Design Report," prepared for the Wireless Communications Bureau of the Federal Communications Commission. http://wireless.fcc.gov/auctions/default.htm?job=papers_studies#experiments.

Goeree, J. K., C. A. Holt, and J. O. Ledyard (2006) "An Experimental Comparison of the FCC's Combinatorial and Non-Combinatorial Simultaneous Multiple Round Auctions," prepared for the Wireless Communications Bureau of the Federal Communications Commission. http://wireless.fcc.gov/auctions/default.htm?job=papers_studies#experiments

Klemperer, P. (2002) "What Really Matters in Auction Design," *Journal of Economic Perspectives*, 16, Winter, 169-189.

Kwasnica, A. M., J. O. Ledyard, D. Porter, and C. DeMartini (2005) "A New and Improved Design for Multi-Object Iterative Auctions," *Management Science*, 51(3), March, 419-434.

Ledyard, J., D. Porter, and A. Rangel (1997) "Experiments Testing Multiobject Allocation Mechanisms," *Journal of Economics and Management Strategy*, Vol. 6(3), 639-675.

Li and Plott (2005) "Tacit Collusion in Auctions and Conditions for Facilitation and Prevention: Equilibrium Selection in Laboratory Experimental Markets," Caltech Working Paper 1223.

McAfee, P. R. and McMillan, J. (1996) "Analyzing the Airwaves Auction," *Journal of Economic Perspectives*, 10(1), 159-175.

McCabe, K., Rassenti, S. and Smith, V. (1989) "Designing 'Smart' Computer Assisted Markets," *European Journal of Political Economy*, 5, 259-283.

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

Porter, D. P. (1999) "The Effect of Bid Withdrawal in a Multi-Object Auction," *Review of Economic Design*, 4(1), 73-97.

Porter, D., S. Rassenti, A. Roopnarine and V. Smith (2003) "Combinatorial Auction Design," *Proceedings of the National Academy of Sciences*, 100(19), September, 11153-11157.

Rassenti, S., Smith, V. and Bulfin, R. (1982) "A Combinatorial Auction Mechanism for Airport Time Slot Allocation," *Bell Journal of Economics*, 13, 402-417.

Appendix A: Rules for the SMR Laboratory Auctions

Rounds and Bid Structure: All licenses are put up for bid simultaneously, and participants may only submit bids on individual licenses. The auction consists of successive rounds in which participants may place bids. Following each round, the high bid for each license is posted. These high bids then become the standing bids for the subsequent round.

Acceptable Bids: In the first round, an acceptable bid must be equal to or exceed the initial price of 0 by 5 points (each point equaled 40 cents in the experiment). Subsequently, in order to be acceptable, a bid must exceed the provisionally winning bid for the license by at least 5 points. Bidders are given the choice of making one of eight incrementally higher bids (in 5 point increments).

Bid Withdrawal: Each bidder has at most 2 rounds in which they are permitted to withdraw any of their provisionally winning bids. After the withdrawal, the seller becomes the provisionally winning bidder for the withdrawn license and the minimum acceptable bid in the following round equals the second highest bid received on the license, which may be less than or equal to (in the case of tied bids) the amount of the withdrawn bid. A withdrawing bidder pays a penalty equal to the maximum of zero or the difference between the price at which the bidder withdrew its bid and the final sale price in the current auction. If the license goes unsold, the bidder would normally be responsible for paying the difference between the withdrawn bid and the sale price in a subsequent auction, plus a small percentage penalty of 3%. In the experiment, there is no subsequent auction, so these penalties for the case of an unsold license were implemented by requiring that the bidder pay a penalty of 25% of the withdrawn bid.

Bidding Eligibility and Activity: Each license in the experiment is assigned one bidding unit. The total number of bidding units available to the bidder establishes the bidder's maximum "eligibility" to bid. National bidders begin each auction with 6 activity units and regional bidders begin with 3. In each round, a bidder's activity is calculated as the number of licenses for which that bidder is a provisional winner, plus the number of licenses for which acceptable bids are submitted. If a bidder's activity falls below the bidder's current activity limit, that limit is reduced to equal the bidder's actual activity. There were no activity rule wavers in the experiment, so a reduction in activity would put an upper limit on the bidder's activity for all subsequent rounds of that auction.

End of Round Feedback: At the end of each round, bidders receive information on all provisionally winning bids, withdrawn bids, and the corresponding bidder ID numbers. Bidders also see the sum of their own values for the licenses that they are provisionally winning and prices that would be paid for the licenses if the auction had ended.

Closing Rule: The auction closes after any round in which no new bids were placed and no bids were withdrawn. In this case provisionally winning bids become winning bids that are used to calculate auction earnings. The experiment did not allow for defaults on payments, so gains were added to cumulative earnings and losses were subtracted.

Appendix B: Rules for RAD (in Bold) and FCC (in Italics)

Rounds and Bid Structure: This is a simultaneous, multi-round auction in which participants may submit bids on individual licenses or on combinations of licenses (packages). Provisionally winning bids are calculated by maximizing seller revenue for the round. (*FCC: Bids have an exclusive OR (XOR) structure in the sense that each bidder can have at most one provisionally winning bid.*)

Acceptable Bids: In the first round, an acceptable bid must be equal to or exceed the minimum opening bid of 0 by 5 points for each license, or by 5 points times the number of licenses in a package. After each subsequent round, "prices" are calculated for each license on the basis of bids received in the previous round. The pricing rule, as specified in Appendix D of the Experiment Design Report, calculates prices that reflect (as closely as possible) the marginal sales revenue of each license based on bids received. Prices for packages are the given by the sum of the prices for each license in the package. In order to be acceptable, a bid must exceed the price of a license or package at least 5 points times the number of licenses covered by the bid. Bidders are given the choice of making one of eight incrementally higher bids (in 5 point increments).

Bidding Eligibility and Activity: Each license in the experiment is assigned one bidding unit. The total number of bidding units available to the bidder establishes the bidder's maximum "eligibility" to bid (3 for regional bidders and 6 for national bidders). **(RAD: In each round, a bidder's activity is calculated as the number of different licenses for which that bidder is a provisional winner or for which that bidder places a bid, either singly or as part of a package.)** (*FCC: In each round, a bidder's activity is calculated as the maximum of (1) the size of the largest package the bidder is provisionally winning and (2) the size of the largest package the bidder is bidding for.*) If a bidder's activity falls below the bidder's current activity limit, that limit is reduced to equal the bidder's actual activity. There were no activity rule wavers in the experiment, so a reduction in activity would put an upper limit on the bidder's activity for all subsequent rounds of that auction.

End of Round Feedback: At the end of each round, bidders receive information on all provisionally winning bids (for licenses and packages) and the corresponding bidder ID numbers. Bidders also see the prices for all licenses, the sum of their own values for the licenses and packages that they are provisionally winning, and the sum of prices that would be paid for those licenses and packages if the auction had ended.

Closing Rule: The auction closes after any round in which no new bids were placed. In this case provisionally winning bids become winning bids that are used to calculate auction earnings. The experiment did not allow for defaults on payments, so gains were added to cumulative earnings and losses were subtracted.

Appendix C: Rules for the Combinatorial Clock (CC)

Rounds and Bid Structure: This is a simultaneous, multi-round auction in which participants may submit bids on individual licenses or on combinations of licenses (packages). Submitted bids stay active until they are removed (just as provisionally winning bids in the other formats are automatically renewed).

Acceptable Bids: In the first round, an acceptable bid must be equal to or exceed the minimum opening bid of 0 by 5 points for each license, or by 5 points times the number of licenses in a package. After each subsequent round, "prices" are calculated for each license on the basis of bids received in the previous round. If a license is in the contained in a bid made by more than one bidder (individually or as part of a package), then the price for that license will rise by the bid increment (5), otherwise the price does not change. For example, if one bidder is bidding on A and AB, and if the only other bidder is bidding on B, then the price of B will increase and the price of A will stay unchanged. Prices for packages are the given by the sum of the prices for each license in the package, so the price of a package can increase by at most the product of the bid increment and the number of licenses in the package.

Bidding Eligibility and Activity: Each license in the experiment is assigned one bidding unit. The total number of bidding units available to the bidder establishes the bidder's maximum "eligibility" to bid (3 for regional bidders and 6 for national bidders). In each round, a bidder's activity is calculated as the number of different licenses for which that bidder is a provisional winner or for which that bidder places a bid, either singly or as part of a package. For example, a regional bidder with an initial activity limit of 3 who is provisionally winning license A would be able to bid on packages BC and on ABC, but not on BC and BE. If a bidder's activity falls below the bidder's current activity limit, that limit is reduced to equal the bidder's actual activity. There were no activity rule wavers in the experiment, so a reduction in activity would put an upper limit on the bidder's activity for all subsequent rounds of that auction.

End of Round Feedback: At the end of each round, bidders receive information on all prices and submitted bids (for licenses and packages), with the corresponding bidder ID numbers.

Closing Rule: The auction generally closes after any round in which there is no excess demand, i.e. no license is in the "bidding basket" of more than one person. However, if there is excess supply at this point (one or more unclaimed licenses), then a revenue maximization routine is run using all submitted bids in all rounds in order to arrange the sale of all licenses. If the resulting allocation displaces the sole remaining bidder for any of the licenses, the auction is restarted and the clock prices on those licenses are raised to let those bidders have the chance to reclaim them.

Appendix D: Efficiency, Revenue, and Profits by Auction Format

	HOHC				LOHC				HOLC				LOLC			
	SMR	CC	RAD	FCC	SMR	CC	RAD	FCC	SMR	CC	RAD	FCC	SMR	CC	RAD	FCC
Efficiency																
Wave 1	86%	92%	89%	88%	86%	89%	93%	89%	96%	92%	95%	88%	96%	82%	94%	93%
Wave 2	85%	94%	89%	93%	79%	91%	89%	91%	96%	93%	95%	84%	97%	93%	93%	86%
Wave 3	84%	86%	94%	93%	85%	87%	94%	87%	98%	96%	98%	93%	96%	96%	99%	92%
Average	**85%**	**91%**	**90%**	**91%**	**83%**	**89%**	**92%**	**89%**	**97%**	**94%**	**96%**	**88%**	**96%**	**90%**	**95%**	**90%**
Revenue																
Wave 1	54%	70%	56%	46%	40%	59%	50%	49%	30%	46%	40%	21%	30%	47%	29%	38%
Wave 2	55%	71%	52%	44%	42%	56%	38%	48%	32%	42%	33%	15%	28%	40%	41%	36%
Wave 3	42%	51%	43%	41%	38%	45%	42%	33%	30%	35%	23%	28%	24%	39%	34%	23%
Average	**50%**	**64%**	**50%**	**44%**	**40%**	**53%**	**43%**	**43%**	**31%**	**41%**	**32%**	**21%**	**27%**	**42%**	**35%**	**32%**
Profit Nationals																
Wave 1	3%	3%	10%	10%	4%	1%	1%	-3%	4%	2%	-3%	5%	3%	2%	0%	2%
Wave 2	-4%	3%	6%	7%	-2%	3%	4%	2%	4%	2%	3%	1%	2%	1%	4%	4%
Wave 3	1%	5%	9%	6%	1%	2%	1%	2%	3%	0%	3%	4%	1%	1%	4%	2%
Average	**0%**	**4%**	**8%**	**8%**	**1%**	**2%**	**2%**	**0%**	**3%**	**2%**	**1%**	**3%**	**2%**	**1%**	**3%**	**3%**
Losses	2%	0%	0%	0%	0%	0%	0%	2%	0%	0%	1%	0%	0%	0%	0%	0%
Profit Regionals																
Wave 1	29%	19%	23%	32%	48%	29%	41%	43%	62%	44%	58%	63%	63%	32%	64%	53%
Wave 2	33%	19%	31%	42%	38%	32%	47%	42%	60%	49%	59%	68%	66%	52%	48%	46%
Wave 3	41%	31%	41%	46%	46%	40%	51%	52%	65%	60%	72%	60%	71%	56%	61%	67%
Average	**34%**	**23%**	**32%**	**40%**	**44%**	**33%**	**47%**	**46%**	**63%**	**51%**	**63%**	**64%**	**67%**	**47%**	**58%**	**55%**
Losses	0%	0%	0%	0%	0%	0%	1%	0%	0%	0%	0%	0%	0%	1%	1%	1%
# Rounds	**15.2**	**22.8**	**24.2**	**17.1**	**16.3**	**22.8**	**22.2**	**18.7**	**13.7**	**18.1**	**16.4**	**17.9**	**14.9**	**18.8**	**15.8**	**21.8**
Unsold Licenses	**4%**	**1%**	**0%**	**0%**	**7%**	**0%**	**0%**	**0%**	**0%**	**1%**	**0%**	**0%**	**0%**	**0%**	**0%**	**0%**