

Winning Play in Spectrum Auctions*

Jeremy Bulow, Jonathan Levin, and Paul Milgrom[†]

February 2009

Abstract

We describe factors that make bidding in large spectrum auctions complex — including exposure and budget problems, the role of timing within an ascending auction, and the possibilities for price forecasting — and how economic and game-theoretic analysis can assist bidders in overcoming these problems. We illustrate with the case of the FCC's Advanced Wireless Service auction, in which a new entrant, SpectrumCo, faced all these problems yet managed to purchase nationwide coverage at a discount of roughly a third relative to the prices paid by its incumbent competitors in the same auction, saving more than a billion dollars.

*The authors advised SpectrumCo in FCC Auction 66. We thank John Hegeman and Marcel Pribsch for excellent research assistance on this paper. Levin acknowledges research support from the Toulouse Network on Information Technology.

[†]Stanford University Graduate School of Business (Bulow) and Department of Economics (Levin and Milgrom), and NBER (Levin); email: jbulow@stanford.edu; jdlevin@stanford.edu; pmilgrom@stanford.edu.

1 Introduction

Since being pioneered by the U.S. in 1994, simultaneous ascending auctions have become a common mechanism to allocate spectrum rights.¹ Spectrum auctions can involve billions of dollars and companies bidding in these auctions regularly create specialized bidding teams and hire experts to develop bidding strategies. Nevertheless, the results can be surprising. In the FCC's auction of Advanced Wireless Service spectrum, price arbitrage failed so dramatically that one new entrant was able to purchase essentially nationwide coverage for about a third (more than a billion dollars) less than what incumbent carriers paid for equivalent spectrum in the same auction (Table 1). At the same time, the other prospective nationwide entrant exited the auction early and filed a letter with the FCC claiming that the auction rules disadvantaged new entrants!

Results of this sort raise questions for economists. Does the apparent failure of the Law of One Price indicate a fundamental flaw in auction design? If not, why must such auctions be complicated? What are the issues that create strategic complexity for bidders? And to what extent can the tools of economic theory provide insights that facilitate effective bidding in highly complex environments?

We start by explaining some of the reasons why large spectrum auctions are necessarily complicated, and why the Law of One Price can fail so dramatically. We emphasize two difficulties facing bidders: *exposure problems*, which occur when bidders wish to acquire complementary licenses, and *budget constraints*, which we argue are ubiquitous. We explain why these difficulties make bidding in simultaneous ascending auctions complicated, and also why they would complicate bidding in other auction designs.

Exposure problems create fundamental difficulties for a new entrant seeking to compete head-to-head with incumbent nationwide wireless carriers in the US. Such an entrant needs to acquire adequate bandwidth in every major metropolitan area, but because licenses covering cities or regions are sold individually, the entrant could commit to spending billions of dollars on certain spectrum licenses before discovering that the total price for the bundle of licenses it seeks makes the whole entry unaffordable or unprofitable. It could then be left to dispose

¹Among the countries that have used such auctions to allocate spectrum are the US, Canada, Mexico, Australia, New Zealand, India, Germany, the Netherlands and the United Kingdom.

of extensive holdings at fire-sale prices.²

The exposure problem, as well as the difficulties created by budget constraints, arise because there is uncertainty about the final auction prices. A bidder who knew what final prices would be would face no exposure problem and have no difficulty deciding how to allocate its limited budget. Information early in the auction about the final auction prices can therefore be extremely valuable to bidders.

Remarkably, it turns out that in large spectrum auctions, information sufficient to forecast final price levels is often available early in the auction. We document this previously unnoticed pattern using data from large FCC auctions. We also provide a simple theory that is broadly consistent with the facts. According to our theory, it is bidders' *budgets*, as opposed to their *license values*, that determine average prices in a spectrum auction.

We then explore the dynamics of simultaneous ascending auctions, in which prices of various licenses may follow a variety of increasing paths. We show that bidders facing exposure problems and budget constraints may wish to manipulate the price paths so larger licenses reach their final prices earlier in the auction than smaller ones. We describe the tactics available to bidders to accomplish that. And we explain the sometimes conflicting interests of new entrants regarding auction timing. Finally, we explain how bidders facing competitors who must deal with exposure and budget problems can disadvantage them by manipulating the price path in other ways.

We illustrate the practical application of these ideas using the experience of the U.S. Advanced Wireless Service auction mentioned above. In that auction, held in the late summer of 2006, the FCC auctioned 90 MHz of nationwide spectrum divided into 1122 licenses. The sale, in which winning bids ultimately totalled 13.9 billion dollars, attracted 168 bidders including two potential nationwide entrants: a consortium of cable television companies and a rival consortium of satellite television companies. Because of budget and exposure problems the potential entrants faced by far the most difficult strategic decisions.

During the auction, the satellite consortium exited earlier than any other major bidder, without buying a single license. The cable consortium, bidding under the name SpectrumCo,

²This “exposure problem” has been the driving force behind attempts to create new auctions in which bidders can bid for packages of licenses. Essays in the book by Cramton, Shoham, and Steinberg (2005) deal with various aspects of this mechanism design problem.

acquired licenses covering 91.2 percent of the U.S. population at prices that were much lower than those paid by the other large buyers. At the per unit prices paid by the major incumbent carriers, SpectrumCo’s licenses would have cost more than 3.5 billion dollars — as it worked out, SpectrumCo paid less than 2.4 billion dollars.³ While luck surely contributed to this outcome, we describe in the final section of the paper how the elements of auction strategy analyzed here, specifically tactics to control the pace of the auction and decisions guided by budget-based price forecasts, put SpectrumCo in a position to be lucky.

2 Why Spectrum Auctions Are Complicated

Spectrum auctions in the United States and other countries have typically been conducted using a simultaneous multiple round format. There are both practical and theoretical rationales for this choice. For instance, under certain conditions a simultaneous ascending auction results in competitive market-clearing prices. Suppose that bidders view licenses as substitutes and that bidding is “straightforward,” meaning that every round each bidder makes offers on the set of licenses that give it the most surplus at current prices. Milgrom (2000) proved that prices will then rise and ultimately stop at approximate competitive equilibrium levels (see also Gul and Stacchetti, 2000).

As an example, consider an auction like the British 3G auction held in 1999. Five national licenses were sold, with no bidder being allowed to buy more than one. Two licenses were bigger than the other three. In this auction, straightforward bidding, with each bidder making the qualifying offer that would give it the most surplus if the auction ended immediately, was a sensible strategy. The simple design helped allocate the spectrum to those willing to pay the most.

In auctions where bidders can buy multiple items, however, both bidders and auction designers face more serious challenges. For starters, to keep the auction moving forward the FCC employs an “activity rule” requiring each bidder to make offers in each round on at least

³The most common unit for measuring the size of a spectrum license or a collection of licenses is MHz-pop, which is calculated by multiplying the bandwidth in MHz by the population covered. Due to both spectrum scarcity and differential build-out costs, prices of large urban areas are higher in MHz-pop, so our calculation may somewhat understate SpectrumCo’s price differential.

a certain percentage of the “quantity” of spectrum which it is eligible to buy, the percentage increasing across two or three auction stages. If a bidder makes offers on a smaller amount of spectrum, its eligibility is reduced.

While the reasons for such an activity rule are compelling, the rule makes it more difficult for bidders to move back and forth between say a 30 MHz license and a 10 MHz license covering the same geographic area but only absorbing a third as much eligibility. Even moving back and forth between a 30 MHz license and three 10 MHz licenses can be tricky, as doing so may require being outbid on the three smaller licenses simultaneously. This barrier to arbitrage helps create the possibility of large price differences among nearly identical spectrum footprints.

Activity rules are one reason that in spectrum auctions bidding activity often starts on the larger licenses and then moves to the smaller licenses (though there are also strategic reasons for this, as we will explain below). Figure 1 illustrates the general pattern for the AWS auction. The two curves plot the round-by-round fraction of bids, by number, that are made on licenses that are larger or smaller than the median license, according to the FCC’s quantity measure (“points”). In this auction, the larger licenses on average saw more bidding for the first 50 or so rounds, after which the pattern reversed.⁴

Of course, there is much more variation in the underlying data. Some small markets clear early, and some large markets can clear quite late. To provide a sense of the variability, Figure 2 plots the round of the last bid for each license against the size of the license for both the AWS auction and FCC Auction 35, another large auction. What is particularly important to note here is the wide range of rounds in which different licenses received their final bids. It is exactly this dispersion that creates problems for a bidder trying to assemble a package and worried about a failed aggregation. We will show below that it also gives these bidders an incentive to manage the pace of the auction, rather than bidding in straightforward fashion.

Our model will assume that a bidder who needs two licenses and acquires only one cannot resell its license. That assumption requires some justification. After all, the winning bidder could try after the auction to sell to the bidder that placed the final competing bid on the

⁴The figure ends after Round 82, at which point there were 75 bids. In later rounds, there were never more than 50 bids and generally far fewer, and the fractions plotted in the chart fluctuate widely due to the small number of bids.

unwanted license, often only one increment below the sale price. The losing bidder, however, may have redirected its limited budget elsewhere, or met its needs some other way, or it may also have attempted and failed to assemble a collection of licenses, eliminating its interest in the single license. Furthermore, even if there is continued interest, the underbidder will know that it offered the highest alternative price, putting it in a strong bargaining position. So while unwanted licenses do often have some salvage value, rational bidders anticipate incurring significant losses in trying to re-sell them.

A second important complication for bidders who cannot forecast final prices arises from limited budgets. Suppose a bidder is targeting two particular licenses that it believes would be profitable at prices up to 200 million and 100 million dollars, respectively, but has been able to raise only 150 million dollars in capital. If the rules require that the bidder remain active on both licenses or else lose the eligibility to purchase both, and the current price of each license is 50 million, what should the bidder do? If it continues bidding on only the more valuable license, it passes up the opportunity to win both. But if it bids for both, it might win the less valuable license in the current round and later find that its budget constraint blocks it from buying the other license, which may be a much better bargain.

Theoretically, some of these difficulties could be addressed using alternative auction designs. For instance, a natural way to address the exposure problem is to permit package bids as well as bids on individual licenses.⁵ Package bidding, however, comes with its own difficulties including complexity problems (there may be very many potential packages!), coordination problems (bidders need to make bids that fit together in reasonable ways) and strategic problems which depend on the particular auction rules.

To illustrate some of the issues, consider the Vickrey auction. In a Vickrey auction, each bidder can make offers on any license packages, and licenses are allocated to maximize the total bidder value assuming that bids reflect true value. Each bidder receives a surplus equal to the difference between the maximized value conditional on its participation and absent its participation. So the price a bidder pays equals its stated value for the licenses it receives, less its incremental contribution to auction surplus. Under well known assumptions

⁵The FCC's recent auction of 700 MHz spectrum allowed a very limited set of package bids. The upcoming British auction of WiMax spectrum will use an even more ambitious package bidding design.

bidders have an incentive to bid straightforwardly in the Vickrey auction and the allocation is efficient.

Unfortunately the Vickrey auction suffers from very serious problems, including low revenue and vulnerability to collusion. Say a package bidder offers 10 for licenses A and B (but nothing for the individual licenses). If two individual bidders offered 9 for each of the individual licenses then each would win at a price of 1! If the bidders each had true values of 6 this would harm revenues; if they had true values of 2 then efficiency would be harmed as well. While in most markets collusion is difficult because it requires the cooperation of almost all competitors and particularly of the highest bidders, in Vickrey-Clark-Groves mechanisms it is not unusual that a successful collusion can be carried out by just two bidders.⁶

Budget constraints also create problems for a Vickrey auction. Consider the bidder above that valued license A at 200 million and license B at 100 million, and had a budget of 150 million. If its bids report its maximum price for each license or bundle of licenses, it will offer 150 for A alone, 100 for B alone, and 150 for the pair. The mechanism will treat it exactly the same as a bidder who places *no value* on B in the event it receives A, so that the bidder will never win both licenses.⁷ Evidently, bidding its maximum price can be a very poor strategy for a budget-constrained bidder in a Vickrey auction.

Finally, a drawback of Vickrey auctions and other “one-shot” auctions is that they limit the ability of bidders to learn from the bids of others. For example, a bidder may think that it is the most efficient operator in market A and therefore “should” win a license there. It may also think that it probably “should” lose out in market B. Finally, it may believe that a license in market B will be worth exactly as much to it as a license in market A. However, it has an extremely noisy estimate of what that value is. How should this bidder proceed in a Vickrey auction? Should it bid the same amount for each license, believing they are equally valuable? What if its value estimates are much higher than most of its competitors? Will it overspend and buy licenses where it is not the efficient operator? On the other hand, what if its estimates are too low? Will it be shut out of licenses that it knows it should win?

⁶This example was presented by Jeremy Bulow at the FCC’s May 2000 conference on combinatorial bidding. More complete descriptions of the faults of the Vickrey auction can be found in Ausubel and Milgrom (2005) and Rothkopf (2007).

⁷Hegeman (2008) studies the possibility of designing an efficient auction for budget constrained bidders.

These problems have led many economists concerned with auction design to favor multiple round auction formats when there is substantial price and value uncertainty.⁸ Given this, we now take up how bidders in simultaneous multiple round auctions can deal with the strategic complexities, particularly the exposure and budget problems.

3 Price Forecasts and Bidder Budgets

3.1 The Price Forecasting Surprise

For a bidder facing a serious exposure problem, such as SpectrumCo in the AWS auction, the central strategic question is whether and when to exit the auction. An accurate early prediction of final prices can allow the bidder to avoid both kinds of exit mistakes, namely, the mistake of exiting too early, when final prices turn out to be low enough for successful entry, and the mistake of exiting too late, when final prices are found to be too high only after the bidder has won some licenses. Accurate price prediction is also valuable for bidders without an exposure problem but facing serious budget constraints, because it allows them to focus their spending on the licenses that will prove to be the best values.

Spectrum auction prices, however, can be hard to predict before the auction. The spectrum offered in each auction often has its own unique characteristics or restrictions imposed by the regulator, and even when an attempt is made to control for these differences, spectrum prices have fluctuated wildly over time. Table 2 shows the average prices per MHz-pop for 9 FCC spectrum auctions between 1995 and 2008. The variation is dramatic and much of it not easily explained by the nature of the spectrum for sale or industry conditions. Even forecasts made just prior to an auction by investment banks tend to have high variance. Prior to the AWS auction, analyst estimates of auction revenue ranged from \$7 to \$15 billion. For the recent 700 MHz auction, they varied over an enormous range – from \$10 to \$30 billion.⁹

⁸There are also important considerations, such as avoiding collusion or encouraging entry, that can weigh against multiple round formats or at least in favor of a shorter number of rounds (Klemperer, 2002). Also, because our focus here is on spectrum auctions run by governments we assume that the seller has a great deal of power in setting the rules. In other situations, buyers may be able to take actions that disrupt or pre-empt an auction, or may refuse to participate unless they deem the rules sufficiently favorable to their interests.

⁹In both of these cases, the final auction revenue fell between the extremes. In other cases, such as in some of the European 3G auctions, revenues have greatly exceeded or fallen short of analyst expectations.

Can information from early bidding usefully reduce the uncertainty? In an ascending auction for a single item, bidders who quit early provide some statistical information, but the only certain conclusion when the price reaches p is that the final price will be no lower. With multiple items being sold, more information may be available. High prices per MHz-pop for licenses that have seen early bidding can be a clue about the values bidders place on licenses that have not yet seen much activity. In fact, we have found that an even simpler approach proves surprisingly powerful for forecasting prices and guiding bidding decisions. This approach focuses attention not on bidder values, which are so emphasized in traditional auction theory, but on bidder budgets.

In major spectrum auctions, even large corporations need to raise or put aside money in advance to finance their spectrum purchases. Many of these companies also have a broad set of target licenses. If these licenses are substitutes and the budget constraint is binding, the bidder's optimal purchase will involve spending its whole budget or nearly so. Of course not every bidder falls into this category. For bidders with tight budgets and narrow interests, or for entrants with all-or-nothing goals, rising prices could lead them to spend zero once the prices of target licenses rise too high.

If bidders in the first category account for enough of the money in the auction, a previously unexplored pattern becomes identifiable in the data. Define a bidder's *exposure* to be the sum of all of its bids in a given round, including its standing high bids from the prior round and all of its new bids in the current round, whether provisionally winning or not. This is the largest amount that a bidder might have to pay if all of its bids were to become winning. If a bidder faces a binding budget constraint and has broad interests, then as prices increase from round to round, its total exposure will eventually level off at an amount approximating its budget. If all bidders were to fall in this category, then the total exposure of all bidders in the auction would rise to the level of the aggregate bidder budgets and level off, forecasting the final auction prices. As prices rise, bidders will narrow the set of licenses on which they bid, the identities of the provisionally winning bidders on various licenses will change, and total winning bids will continue to rise, but final total winning bids will be forecast early and well by total exposure.

The same idea can be expressed using a supply-demand analysis. Assume that bidders

are required to make offers each round on about as much spectrum as they are interested in buying. Then we have

$$\text{Exposure} = \text{Total Price} \times \frac{\text{Demand}}{\text{Supply}}$$

For example, if at current prices all the licenses sell for \$10 billion in the auction and the potential winning bids or exposure in a round is \$15 billion, then at the current prices the (price-adjusted) quantity of spectrum demanded is 1.5 times supply. The equilibrium total price, where demand equals supply, will be determined by the aggregate elasticity of demand. The budget hypothesis tells us that by this stage of the auction many bidders will be constrained by their budgets and so have an (uncompensated) elasticity of demand of -1 .¹⁰

In this example, say that when the total price was \$15 billion demand was composed of budget bidders offering \$10 billion and other bidders offering \$5 billion. Assume that the non-budget bidders exhibit constant elasticity of demand in aggregate. Because the budget bidders anchor elasticity at -1 , revenue will be relatively close to the exposure of \$15 billion for a wide range of elasticity. If for example the non-budget bidders have an elasticity between -2.19 and $-.48$, final revenue will be between \$13.5 billion and \$16.5 billion, or within 10 percent of exposure. By contrast, if there were no budget bidding and the same range of possibility of elasticity, then final prices could only be estimated within a factor of two — between \$12.15 billion and \$24.30 billion. Clearly, if the budget hypothesis has validity it is a tremendous aid to forecasting.

Is auction data consistent with the budget hypothesis? Figure 3 shows the pattern of total exposure and total prices in Auction 35, which was the largest US spectrum auction in the years before the AWS auction. At round 10, total revenue in the auction was still less than one-third of their eventual level, but total exposure had approached its final level.

¹⁰Technically, two small adjustments to this formula are needed when applying to real spectrum auctions. First, bidders are typically required to make bids for only 95 percent of the spectrum they might still demand, even in the final stage of an auction, causing the formula to potentially understate demand. Second, some exposure is accounted for by bids that were at a price one increment below the current price, and not all those bids would be renewed at the current price, causing the formula to potentially overstate demand. Because these two adjustments to the supply-demand calculation work in opposite directions they partially cancel one another.

Forecasting that total revenue would be equal to total exposure from that point forward would lead to errors that are mostly less than 10%, close enough to guide some of the most critical strategic calculations. A potential new entrant who decided, based on that forecast, that prices would become too high could stop bidding while prices were still far below their final levels. The entrant would likely be topped on most or all of its provisionally winning bids and, even if it were not, its early withdrawal would mean that it acquired licenses at only a small fraction of the average auction price, greatly reducing any expected loss on resale.

Figure 4 shows the similar pattern of exposure and revenue for the AWS auction. Again, total exposure provides a remarkably accurate early forecast for total prices in the auction. Exposure peaked at \$14.2 billion in round 11 and final auction revenue was \$13.9 billion. The large drop in exposure in round 13 is largely due to the exit of Wireless DBS, a joint venture of the two satellite TV companies Echostar and DirecTV. In Round 12, Wireless DBS's exposure was \$2.025 billion; it dropped to \$196 million in round 13 and subsequently to zero. From round 15 onwards, however, a bidder who estimated final total prices to be equal to current total exposure would never have made an error larger than 10%, despite the fact that the total price was still 40% below its final value.

In the AWS auction, the ability to forecast prices early in the auction had another key implication. Early bidding in that auction focused almost entirely on the 40 MHz of spectrum that was divided into large REAG licenses, before turning to the 50 MHz of spectrum that was divided into smaller EA and CMA licenses. By round 15, it was possible to forecast that cumulative high bids on the REAG licenses were so high relative to the total budgets in the auction that the smaller licenses would sell for a steep discount. This allowed SpectrumCo, alone among the major bidders, to make an early switch to the smaller licenses.

Further evidence on the budget constraint theory can be seen in Figure 5, which plots the exposure of the largest individual bidders in Auction 35 and the AWS auction. These figures suggest more than one pattern of bidding. In Auction 35, the behavior of AT&T, Cingular, and many of the smaller bidders suggests a binding budget that could have been inferred early on. Verizon eventually may have hit a budget constraint, but if so not until relatively late in the auction. Similarly, all of the major winners in the AWS auction —

T-Mobile, Verizon, SpectrumCo, MetroPCS, Cingular, Leap/Denali and Barat — exhibit budget-constrained patterns. Yet there are also clear exceptions. Two large bidders, the Dolans and Wireless DBS, stand out. Wireless DBS was an all-or-nothing entrant. The Dolans were bidding on New York licenses, presumably to complement their New York cable franchise Cablevision. When the prices became prohibitive, these bidders found no desirable substitutes upon which to spend their budgets, so they left the auction.

Using exposure to forecast prices would have worked well in many past FCC auctions, but not in every case. Exposure in smaller auctions sometimes has peaked well above final revenue, as one might expect with independent ascending auctions not tied together by a budget constraint. Figure 6 shows the ratio of maximum auction exposure to final auction revenue for ten previous FCC auctions. For the larger auctions, exposure does not rise much above ten percent over final auction revenue, but there is greater variance for the smaller auctions. We note, however, that in some of these auctions exposure peaked for just one or two rounds. Figure 6 also displays, using smaller hollow squares, the same exposure to final revenue plot using instead the third highest round of auction exposure. From this plot, we see that exposure remained well above final revenue for a significant period in only one of the ten auctions.

What Figure 6 does not show is that in these auctions, as in Auction 35 and the AWS auction, exposure also climbed much faster than revenue and so provided a usefully *early* forecast of final auction revenues. We document this phenomenon in Table 3 which, for each auction, reports (1) the round in which revenue reached 90% of its final level, (2) the round in which exposure reached 90% of final revenue, and (3) the ratio of revenue in that round to final revenue. The choice of 90% is, of course, arbitrary, but it is a reasonable choice because 10% is the smallest bid increment that the FCC used for individual licenses in the AWS auction. Table 3 shows that in nine of ten cases auction exposure reached 90% of final revenue at a point where auction revenue was under half its final level.

To understand why a forecast is valuable for a new nation-wide entrant, suppose the entrant has a target budget, perhaps with 10% upward flexibility. From the evidence above, even a very simple budget-forecast strategy, such as “exit the auction only when the forecasted price for a national footprint based on the total exposure price exceeds the target

budget,” would have performed extremely well. It rarely would have recommended an inappropriate exit, and when exit was recommended, prices would have been quite low relative to the auction close.

Of course, we are not advocating a mechanical approach to bidding, and the results of the recent FCC Auction 73 (for 700 MHz spectrum) provide one reason. In that auction, exposure peaked at \$25.6 billion in round 27, but the final auction revenue was just \$19.6 billion, suggesting a dramatic failure of the budget theory. The gap, however, was attributable in large part to a single bidder, Google, which had a provisionally winning bid of \$4.7 billion on a national package license through round 27, but exited the auction when its bid was topped by Verizon. The prospect of exactly this behavior had been widely discussed even before the auction, because of Google’s unusual role and objectives.¹¹

3.2 Why Bidder Budgets?

Why do the teams representing large bidders in spectrum auctions face budget constraints? Superficially, the answer appears simple. Bidding in a spectrum auction requires a substantial amount of cash-on-hand, and raising this money from external capital markets takes time. In turn the capital markets may want to deliver money against a promised acquisition. If this pattern is optimal, then it is hardly surprising that the same pattern of capital budgeting could emerge in companies funding a division bidding in a spectrum auction.

The harder questions concern why this pattern of capital budgeting prevails and why prices vary so widely over time. In practice, incumbent firms can often substitute for additional spectrum by using existing spectrum more intensively, by building more cell sites or by using other spectrum enhancing technologies. It seems unlikely that the shadow cost of spectrum fluctuates so substantially over time. Nevertheless, evidence from behavior in spectrum auctions suggests that bidding teams often face budget constraints and yet have considerable freedom in deciding which licenses to buy within their fixed budgets. Such a

¹¹ Google had lobbied the FCC to include an “open access” band in the auction. Under the auction rules, if the FCC-set \$4.6 billion reserve were met, the winner of that band would be required to allow the operation of devices and software from independent providers (such as Google). If the reserve were not met, then the open-access provision would be removed and the licenses made available for re-auction. Google participated in the auction until the reserve was met and the open-access provision was triggered. It then immediately ceased bidding.

pattern might be rationalized if the bidding team has better information about the relative values of different licenses but also has either different incentives or different beliefs about factors like demand growth that affect the value of the entire business.¹²

4 Controlling the Pace of the Auction

By far the most dramatic moment in the AWS auction occurred in round 9, when SpectrumCo made a jump bid doubling the prices of the large REAG licenses covering the Northeast and Western United States. This move, which we referred to at the time as the “shake-out tactic,” was intended to resolve competitive uncertainty and favorably align relative prices in the auction. In particular, it aimed to alleviate the risk that SpectrumCo might end up purchasing the licenses across the interior U.S. but fail to purchase licenses covering the large cities on the coasts.

In this section, we explain why bidders facing exposure or budget problems almost always have an incentive to control the relative pace of price increases of different items in a simultaneous ascending auction.¹³ In the process, we identify optimal bidding patterns in a stylized model of the auction, characterize the welfare effects, and then explain the practical implications.

4.1 An Illustrative Example

Consider a bidder who is interested in acquiring two licenses, A and B. It is willing to pay 10 for the package but regards each individual license separately as worthless. It does not know the values that others place on the licenses, but thinks that the amount it will have to pay to win license A is uniformly distributed between 0 and 10, and the amount to win B is independent and uniformly distributed between 0 and 6. No package bidding is possible.

¹²Our own experience in these auctions indicates that one difficulty in relying solely on net present value estimates of license values is that these kinds of estimates are extremely sensitive to assumptions about interest rates, demand growth, market share, and so on. For example, we know of a successful bidder that prior to a major FCC auction estimated the value of a Chicago area license at \$30 per covered person, plus or minus \$60. With such a wide range of values, a binding budget constraint may be a sensible way to focus a bidding team on relative values.

¹³Several papers explain why bidders may have a signalling motive to use a jump bid (e.g. Avery, 1998; Brusco and Lopomo, 2002; Zheng, 2006). The rationale we describe is quite different.

Participating in the auction is profitable in expectation — even a brute force strategy of buying both licenses regardless of the price earns an expected profit of 2 — but exactly how profitable depends on the bidding dynamics.

Suppose for example that license B sells first. If the buyer purchases B for p_B , it will certainly want to buy A, but could lose money overall if the price of A goes above $10 - p_B$. Its best strategy is to bid for B only until the price reaches 5, at which point it would make zero in expectation from winning. This strategy gives expected profit of 2.0833. It is better for the bidder if A sells first. Its optimal strategy is then to bid for A until its price reaches 7, allowing it an ex ante expected profit of 2.45.

In a simultaneous ascending auction the buyer ideally would like to see the price of A rise faster than B until it either wins a license or decides to exit. If no license clears earlier, the buyer would have the price of A reach 6 and the price of B reach 2 at exactly the same time, and then quit. This raises its expected profit to 2.533.

Intuitively, the buyer prefers that prices rise in a way that conveys as much information as possible before it must commit to buy or not. That tends to make the buyer prefer a faster increase in the price of the more uncertain license. In this case, there is initially more uncertainty about the price of license A and, so long as the other bidders remain active, the buyer's best policy is to raise the prices until they reach (6,2). At that price vector, the remaining uncertainty about both license prices is the same and winning either license A at price 6 or license B at price 2 leads to expected profits of zero.

This illustrates a general principle: on the buyer's most preferred price path, she exits at a point where it would get exactly zero expected profit from winning *either* license at its current price, given its conditional expectation about the other license price. This principle, which can be inferred from reasoning about the first-order optimality condition, also applies if the bidder needs to assemble multiple licenses, or if the licenses have some stand-alone value.¹⁴

¹⁴While we will focus on controlling prices to manage the exposure problem, we note that a bidder with additive license values may want to do the same if it has a budget constraint. To illustrate, consider a bidder that values license 1 at $2v$ and license 2 at v , but has a budget b with $v < b < 3v$. Ideally, this bidder would like the price of license 1 to rise to $(b + v)/2$ by the time the price on the second license reaches $(b - v)/2$. It would be happy to buy either license at a lower price, regardless of what would be required to buy the other. Once the target prices are reached, however, the buyer cannot afford both licenses, and simply wants to maintain a constant price difference of v between the two licenses. That is, the bidder wants to keep the

4.2 Managing the Exposure Problem

We now consider a more general case where an entrant is willing to pay a premium for two licenses over the sum of its individual license valuations. Let the entrant’s value for license i individually be v_i while the value of the package is $v_{12} > v_1 + v_2$. Suppose the entrant has a budget of b , and for simplicity that $b \geq v_{12}$. There is one competitor on each license (we will generalize this later), and they have independent unknown values c_1 and c_2 , drawn from distributions $F_i(\cdot)$ with densities $f_i(\cdot)$. We will assume that these competitors bid “straightforwardly,” that is, each remains active on its license until the license price exceeds its valuation.

We consider a hypothetical setting where the license prices, denoted p_1 and p_2 , rise continuously and the package bidder can choose the price path as a function of the activity of the other bidders. At any price p , the package bidder knows both p and the prior bidding by the competitors; observing the latter is the same as observing $\min(c_i, p_i)$ for $i = 1, 2$. Denote this historical information by $h(p)$. A strategy for the entrant is a pair $\sigma = (P, d)$ where P is a price path which may depend on the drop-out decisions by other bidders and d is a decision rule specifying whether to exit or continue bidding on each license at each price pair depending on which other bidders still remain.

An expected-profit maximizing entrant will continue bidding on each individual license i at least until the price reaches the entrant’s stand-alone value v_i and will never buy i at a price exceeding its maximum marginal value: $v_{12} - v_j$. A graphical approach will help explain the optimal decision rule d^* in more detail. First, suppose that the price pair (p_1, p_2) is reached and the individual bidder on license i then exits the auction. Conditional on that event, the entrant expects a profit of

$$\pi_i(p_1, p_2) = v_i - p_i + Q_j(p_1, p_2). \tag{1}$$

The final term represents the “option value” of continued bidding on license j , having pur-

relative price higher on the more valuable (or larger) license.

chased license i :¹⁵

$$Q_j(p_1, p_2) = \int_{p_j}^{\max\{b-p_i, p_j\}} \max\{0, (v_{12} - v_i - c_j)\} dF_j(c_j | c_j \geq p_j). \quad (2)$$

Suppose that the entrant selects a path of prices passing through the point (p_1, p_2) and plans to stop bidding on both licenses at that point if both competitors are still active. For a given p_j , increasing p_i slightly raises the entrant's profit if $\pi_i(p_1, p_2) > 0$ and reduces it if $\pi_i(p_1, p_2) < 0$, so optimality requires that

$$\pi_1(p_1, p_2) = 0 \quad \text{and} \quad \pi_2(p_1, p_2) = 0. \quad (3)$$

Figure 7 depicts the two curves satisfying equations (3). To understand the picture, observe first that if $p_i < v_i$, then since $Q_j(p_i, p_j) \geq 0$, $\pi_i > 0$, that is, continuing to bid on license i is profitable. If $p_i = v_i$, there is no immediate payoff to winning license i , but nevertheless $\pi_i > 0$ provided that there is a chance of profiting from the addition of license j , i.e. if $p_j < v_{12} - v_i$. Similarly, it may be desirable to purchase license i even if $p_i > v_i$ provided p_j is sufficiently low. In this region, the curve satisfying $\pi_i = 0$ slopes down because π_i is strictly decreasing in both license prices.

The two curves satisfying (3) must cross at some p^* satisfying $v_i < p_i^* < v_{12} - v_j$. The reason is that if $p_i \geq v_{12} - v_j$, then $Q_j(p_i, p_j) = 0$ and so $\pi_i < 0$: buying license i could not possibly lead to positive expected profits. Also, $\pi_i > 0$ for sufficiently small p_i . In Figure 7, the crossing point is unique, which must be the case given the following condition.

(U) If (3) holds at (p_1, p_2) , then $\left| \frac{\partial Q_2 / \partial p_1}{\partial Q_2 / \partial p_2 - 1} \right| > \left| \frac{\partial Q_1 / \partial p_1 - 1}{\partial Q_1 / \partial p_2} \right|$.

Condition (U) states that any crossing point, the curve defined by $\pi_1 = 0$ is strictly steeper than the curve defined by $\pi_2 = 0$. As both curves slope down, this means at most one intersection.

We can now state our first result.

¹⁵This "option value" analogy can be made precise. If the entrant's budget b is sufficiently large, then $Q_j(p_i, p_j) = \mathbb{E}[\max\{0, v_{12} - v_i - c_j\} | c_j > p_j]$, which is the value of a put option on c_j with exercise price $v_{12} - v_i$, conditional on $c_j \geq p_j$.

Proposition 1 *Assume that (U) holds. The optimal strategies for the entrant are characterized as follows: Raise prices (along any path) to the unique price pair $p^* = (p_1^*, p_2^*)$ that solves (3); drop out at p^* if both competitors are still active at that*

Proposition 2 *point; otherwise, if the individual bidder for license i is the first to drop out, continue bidding on license j until $p_j = \min\{b - p_i, v_{12} - v_i\}$.*

Proof. Recall that any strategy in which the entrant plans its initial exit at a price p such that $p_j > v_{12} - v_i$ cannot be optimal. So we can restrict ourselves to strategies that involve an initial exit price vector \hat{p} satisfying $\hat{p}_j \leq v_{12} - v_i$ for $j = 1, 2$.

Conditional on the initial exit being planned for price vector \hat{p} and other decisions made optimally, the path of prices leading to \hat{p} is irrelevant. Why? If both individual bidders have values above \hat{p} , they will not exit along any path to \hat{p} , so the path does not affect payoffs. If one individual bidder, say i , has a value below \hat{p}_i , it will exit along any path to \hat{p} and its exit will make it optimal for the entrant to remain bidding on j at least to $v_{12} - v_i \geq \hat{p}_j$. By the same logic, if both individual bidders have values below \hat{p} , the entrant will win both licenses regardless of the path toward \hat{p} .

Finally, we argue that a price path leading to initial exit at p^* is best. Consider a strategy with initial exit at $\hat{p} \neq p^*$, where $\pi_i < 0$ for some i . As all paths to \hat{p} yield equivalent payoff, consider the path where just prior to reaching \hat{p} , only the price of license i is rising (such a path must exist because $\pi_i < 0$ at \hat{p} , then $\hat{p}_i > 0$). The entrant does better to follow this path and drop out a bit before \hat{p} . Next, consider a strategy with initial exit at prices \hat{p} satisfying $\pi_i > 0$ and $\pi_j \geq 0$. Rather than exit at \hat{p} , the entrant does better to follow a continuation path in which the price rises just on i , which makes strictly higher expected profits than exiting on one or both licenses at \hat{p} . So $\pi_1(\hat{p}) = \pi_2(\hat{p}) = 0$ at any optimum. *Q.E.D.*

What does the result imply about the direction in which the entrant should push prices? A simple case is where the individual bidder valuations have a constant hazard rate, the same on both licenses. In this case, the entrant's optimal exit point satisfies: $p_1^* - v_1 = p_2^* - v_2$. That is, the bidder will manage prices to equalize the exposure risk across the two licenses.

More generally, the entrant tries to limit exposure risk in the following sense. Along an optimal price path, the entrant may purchase a license at an immediate loss (i.e. at a price $p_i > v_i$), but it never makes a purchase that leads to negative conditional expected profit. That is, along an optimal path, so long as the entrant is active on both licenses, we must have $\pi_1, \pi_2 \geq 0$. In fact, this characterizes the optimal choice p^* . For consider any strictly increasing price path that does not pass through p^* . The entrant would always want to continue bidding beyond the point where $\pi_i = 0$ for some i , and $\pi_j > 0$ for some j , risking a small expected loss if bidder i drops first in favor of a larger expected gain if bidder j drops first.

The case with multiple individual bidders is described in the following proposition.

Proposition 3 *Assume that (U) holds and that there is at least one individual bidder for each license and at least two for some license, with all individual values independently and continuously distributed. An optimal strategy for the entrant is to increase the license prices at any rate until the lowest price \hat{p} at which there is just one other remaining bidder for each license. Then: (1) if $\pi_1(\hat{p}), \pi_2(\hat{p}) < 0$, exit; (2) if $\pi_1(\hat{p}), \pi_2(\hat{p}) \geq 0$, follow the strategy of Proposition 1; or (3) if $\pi_i(\hat{p}) \leq 0, \pi_j(\hat{p}) > 0$, increase price p_j either until $\pi_j(\hat{p}_i, p_j) = 0$, in which case exit, or until the remaining j competitor drops out, in which case increase p_i until it reaches $\min(b - p_i, v_{12} - v_j)$.*

Proof. In searching for an optimal strategy, we can restrict attention to strategies where the entrant never makes an initial exit until \hat{p} , i.e. until there is a single competitor on each license. To see this, consider a strategy that after some history calls for the entrant to exit with multiple bidders left on license i , i.e. with $p_i < \hat{p}_i$. Instead, increase p_i up to \hat{p}_i and execute the same exit. This alternative achieves the same payoff.¹⁶

Consider continuation play from \hat{p} . For cases (1) and (2), optimal behavior follows from the proof of Proposition 1 (just as if $\hat{p} = (0, 0)$). The new case is the one with $\pi_i(\hat{p}) \leq 0, \pi_j(\hat{p}) > 0$. Mimicking the proof of Proposition 1 establishes that there is some $\tilde{p} > \hat{p}$ such that any path from \hat{p} through \tilde{p} is part of an optimal strategy. If $\tilde{p}_i > \hat{p}_i$, then

¹⁶This is obvious if the planned exit was on both licenses or just license i . If the planned exit was on just license j , and the continuation strategy from the initial price pair called for an exit on i between p_i and \hat{p}_i , the entrant can exit on i at \hat{p}_i and guarantee the same (zero) payoff.

one optimal path passes through the price pair (\hat{p}_i, \tilde{p}_j) and then continues to \tilde{p} , but that path has a lower payoff than the same path stopped at (\hat{p}_i, \tilde{p}_j) . This contradiction implies that at the optimum, $\tilde{p}_i = \hat{p}_i$. Therefore at the optimum the entrant increases only the price p_j , and up to the point where $\pi_j(\hat{p}_i, p_j) = 0$. If the last j competitor exits as p_j rises, the optimal continuation on license i follows Proposition 1. Q.E.D.

After the AWS auction, the FCC moved to limit the ability of bidders to engage in jump bidding, which can be a key tool for controlling the pace of the auction. Was the FCC right to do that? Does bidder control of the price path as described damage efficiency?

In the preceding model, regardless of the price path, the entrant's total payment for any licenses it acquires is the sum of its competitors' values, that is, the entrant pays the social opportunity cost of any licenses it acquires. Consequently, an entrant that maximizes its own net profits necessarily maximizes the net auction surplus.

Proposition 4 *If the entrant controls the feasible path of prices to maximize its expected profits, then any effective restriction on the entrant's control reduces expected total surplus from the auction.*

This result should be interpreted with some caveats. First, an auctioneer often has criteria other than total surplus, including revenue, future product market competition, and so forth. Even focusing on total surplus, the model itself omits two potentially relevant considerations that may cause bidders to manipulate prices in a way that distorts efficiency.

One is illustrated by the AWS auction, in which there were *two* potential national entrants, SpectrumCo and Wireless DBS, not just one as specified in our model. In general, two entrants with different values will prefer different price paths, and they cannot both maximize auction surplus. Moreover, an entrant who could choose the price path to maximize its own net profit would not internalize the effect of its choice on the profits and entry decisions of the second entrant, so its choice would not generally be efficient.¹⁷

Even in the model with just one entrant, we have treated individual bidders as passive automata and not as strategic players. A strategic individual bidder for license j might seek

¹⁷One case where the presence of multiple entrants does not create problems for efficiency in our model is if there are two entrants and only one individual bidder competing for, say, license 1. If the entrants are not budget constrained, any price path will result in an efficient allocation.

a path of prices to disadvantage the entrant. Suppose, for instance, that the entrant has zero value for an individual license and a value of 2 for the pair, that the individual bidder for license 1 has a value of 1, and that the individual bidder for license 2 has an uncertain value c which is uniformly distributed on $[0, 3]$.

If the price of license 2 climbs first, the entrant will bid up to 1 and will win both licenses if that is efficient. Bidder 2 will win a license only when $c > 1$ and its profit will be $c - 1$. If bidder 2 could force the price of license 1 to rise first, and force its price to rise first to 1, the entrant would find it unprofitable to bid beyond $1/2$, allowing bidder 2 to acquire its desired license at a price of zero. Bidders can have a similar incentive to drive up the price of non-desired licenses even if there are no value complementarities, but some bidders are budget constrained. In that way, bidding up the price of license 1 may reduce competition on license 2.

4.3 Managing Prices in Practice

A difference between our stylized model and a real auction is that in reality prices do not rise continuously, and a buyer cannot perfectly control the pace of the auction. Nevertheless, the FCC rules do give bidders several ways to influence pacing.

(1) *Holding back demand*: In the early rounds of FCC auctions, bidders need not bid on all the spectrum they are eligible to win. For example, early in the \$7 billion “AB” auction of 1995, bidders could maintain their eligibility by making offers on just one third as much spectrum. So bidders could simply defer bidding on many target properties.

(2) *Parking*. To the extent that activity rules do require a bidder to place bids, it can “park” eligibility by bidding on non-target licenses, planning to switch later to the licenses of main interest. This tactic, too, can affect the relative rate of price increase among licenses.

(3) *Jump bidding*: Though seldom used as a strategic tool prior to the AWS auction, jump bids can similarly raise the price on one license ahead of others. A jump bid made early in the auction can be ineffective in altering the relative rates of price increase, because any changes can be undone by competitors’ responses. A jump bid late in the auction entails some risk of overpaying. This makes the timing and analysis of jump bids subtle, but far from impossible.

5 The AWS Auction

The AWS auction provides an illustration of how jump bidding and the strategic considerations analyzed above can play out in practice. The FCC offered for sale 90 MHz of spectrum covering all of the United States and its territories, divided into 1122 licenses. One 20 MHz layer of spectrum (the “A” band) was divided into 734 Cellular Market Area (CMA) licenses. Two other layers, one 20 MHz and the other 10 MHz (the “B” and “C” bands), were carved into 176 Economic Area (EA) licenses. Finally, three spectrum layers of 10, 10 and 20 MHz respectively (the “D”, “E” and “F” bands), were divided into 12 Regional (REAG) licenses.

Traditional thinking about the exposure problem highlights large licenses as offering the most protection, so the large REAG licenses appeared to offer the easiest route to a nationwide aggregation. Of the 12 licenses covering the US and its territories, a footprint covering the contiguous U.S. required just 6 — the Northeast, Southeast, Great Lakes, Mississippi Valley, Central and West regions. The other 6 licenses in each REAG band covered Alaska, Hawaii and various U.S. territories (e.g. American Samoa), and had much lower values. Based on historical prices, the REAG licenses in different regions were not expected to settle at equal prices, either in absolute terms or on a per MHz-pop basis. For example, the price of a license covering the densely populated Northeast would normally be much higher even on a per MHz-pop basis than, say, the price of the Mississippi Valley license.

The auction attracted 168 bidders, including incumbent carriers Verizon, Cingular, T-Mobile, MetroPCS, and Leap Wireless, and the two potential national entrants: SpectrumCo (the cable consortium) and Wireless DBS (the satellite TV consortium). Bidding was expected to be fierce. Prior to the auction, rough analysis based on budgets strongly suggested that at most one of the national entrants would be able to complete a successful aggregation, and also that the entrants potentially had quite similar resources. As a further complication, there was some concern that incumbent carriers might try to deter entry by purchasing large amounts of the spectrum in key markets such as New York and Los Angeles. Successful national entry without these markets was thought to be impossible, partly because the scarcity of spectrum in these highly-populated markets makes it difficult to buy after-market

spectrum or negotiate a roaming agreement.

We described above how early bidding tends to focus on the largest, most valuable licenses. The AWS auction was no exception: initial bidding centered on the REAG licenses. Nevertheless, a somewhat unusual pattern emerged. The FCC had set essentially uniform starting prices for the licenses (with prices measured on a per MHz-pop basis), with the same minimum percentage price increments for all licenses that attracted high activity. With almost all bids being made at the minimum price increment, prices across the REAG licenses remained uniform even as they climbed far from their starting values.¹⁸

Of course the final market clearing prices were unlikely to be anywhere near uniform. If the early pattern persisted, bidding would likely close on the less valuable REAG licenses such as those covering the Mississippi valley and the Mountain states long before the final prices were determined for the coveted licenses covering the northeast and west coast.

From the perspective of a potential entrant, this timing posed a serious danger. SpectrumCo could wind up winning licenses covering the interior U.S., only to find that price of spectrum covering the coasts had become prohibitive. A further concern was that SpectrumCo and Wireless DBS could each win licenses in the interior U.S., virtually guaranteeing that at least one of them would be left with an economically untenable partial footprint. And, with minimum prices rising at 20% per round and rounds occurring every two hours, these dangers were increasingly imminent.

At this point SpectrumCo executed its “shake-out” tactic, submitting a maximal jump bid on all of the Northeast and West REAG licenses. It submitted these bids on the first round of the Monday morning after the auction started, timed to give competitors only a few hours to respond. The jump bids doubled the prices on the Northeast and West licenses (from roughly \$0.20 per MHz-pop to \$0.40), but SpectrumCo did not assign a large probability to these bids becoming winning. While it would have been satisfied with that outcome, the primary goal was to resolve as much uncertainty as possible and align relative prices.

The jump bids left Wireless DBS in a position of having to raise their bids and loss exposure by hundreds of millions of dollars with just two hours notice and in unexpected

¹⁸Small rural licenses had starting prices that were somewhat discounted (\$0.03 per MHz-pop, as opposed to \$0.05 for all the other licenses). The minimum increment rules also meant that licenses receiving different numbers of bids did not rise in lockstep.

circumstances. Moreover, the shift in relative prices may not have favored their business model, as the satellite television subscriber base skews more toward rural areas than that of cable operators. Wireless DBS responded at first by buying time, taking its first waiver, and then by exiting the auction.¹⁹ Meanwhile SpectrumCo's bids were topped and the REAG prices climbed still higher.

SpectrumCo's attention turned next to its own decision about whether to remain active and how to acquire its desired footprint at the lowest possible price. Forecasting based on budget exposure played a central role. At the end of the day on Monday (following Round 12), auction exposure had peaked at \$14.2 billion. Meanwhile prices across the spectrum bands had diverged so remarkably that the high bids on the 40 MHz of REAG licenses, just 44% of the total spectrum in each area, totaled slightly over \$5 billion, while the high bids on the other 50 MHz totaled just \$737 million (see Figure 8). According to the budget model, the REAG licenses appeared likely to wind up over-priced relative to the smaller EA and CMA licenses unless there was a dramatic shift in bidding intensity. In response, SpectrumCo switched its bidding to the smaller licenses in the first round on Tuesday morning.²⁰

By mid-day Wednesday, the budget model indicated that SpectrumCo's switch might be rewarded handsomely. The cumulative high bids on the 40 MHz of REAG licenses reached \$7.6 billion, versus just \$2.3 billion for the 50 MHz of EA and CMA licenses. Given the earlier exposure peak, the budget forecast implied that the smaller licenses were unlikely to sell for more than \$6.7 billion in total, or \$0.47 per MHz-pop compared to the current and still-rising price of \$0.67 per MHz-pop for the REAG licenses.

It is tempting to ask why other bidders, who had access to precisely the same information, did not identify the same opportunity. Is this evidence of "irrationality" in the bidding? Only in the same sense that chess grandmasters are irrational because they change their

¹⁹This sequence of events does not imply that Wireless DBS erred. It faced, on very short notice, the need to make a multibillion dollar investment decision in an unprecedented auction environment with the timing pattern of bids reversed from what it likely expected. The risks may have been magnified because the expensive licenses were not necessarily in the consortium's strongest markets. And, even if its bidding team could sort out the situation in real time, continuing to bid would have required educating and then getting approval from the senior management of the two consortium companies within just hours, certainly a daunting task.

²⁰The switch was not totally committing in the sense that it would have been possible, and at least early on not terribly difficult, to switch back into the REAG band. Nevertheless, there was a strong view at the time that the EA licenses offered the best route to success.

play over time. What the pattern of bidding highlights is that the auction game is similarly complicated, and that the incumbent bidders who dominated the auction, not facing the same challenging entry decision as SpectrumCo, may have devoted less resources to forecasting final prices early in the auction.

After shifting to the smaller licenses, SpectrumCo still faced an exposure problem, this time the problem of putting together within its budget enough of the 176 EA licenses to have a meaningful national footprint. faced another problem in completing a successful aggregation. At the time, the budget theory indicated that the 50 MHz of EA and CMA licenses would sell for a total in the vicinity of \$6.5 billion, meaning that SpectrumCo could likely acquire 20 MHz nationwide, or 40% of the smaller licenses, for roughly \$2.6 billion.²¹ For a new entrant, however, it is the major markets that are essential for a successful aggregation. So to minimize risk, and again control relative prices, SpectrumCo initially bid for 30 MHz in the major areas, eschewing less valuable licenses. SpectrumCo then adopted a “steadfast” posture, strongly defending the 20 MHz B band licenses, while slowly allowing itself to be bid off the 10 MHz C band licenses and using the freed up eligibility to fill in the non-major market B band licenses.²²

This strategy led to some interesting late auction decisions as the budget forecast was borne out. One implication of the budget hypothesis, in its strictest form, is that a bidder that reduces its bidding by a dollar will reduce total prices by a dollar. SpectrumCo was looking to purchase roughly 40% of the EA and CMA spectrum, so when its high bid on a license such as Lexington, KY was topped, conceding the license promised to save roughly 40% of the bid amount in spending on other licenses. At the same time, there was a strategic reputational issue. A show of weakness might encourage another round of bidding on a SpectrumCo license such as New York. With the large bid increments, such a development could easily end up costing \$50 million dollars. Balancing these considerations, SpectrumCo maintained its steadfast posture, almost without exception, to the end of the auction.²³

²¹While the exact size of SpectrumCo’s budget cannot be disclosed directly (and indeed was not known precisely by its outside advisors at this point in the auction), Figure 6 shows its bidding topping out several times around \$2.5 billion.

²²SpectrumCo ultimately purchased just one C block license, covering Dallas.

²³There were only four markets priced at over \$2 million in which Spectrumco did not acquire licenses — St. Louis (\$23.5 million), Cincinnati (\$21.9 million), Greenville, S.C. (\$5.2 million), and Lake Charles, LA (\$3.6 million).

When the auction finally ended, the REAG licenses had sold for an average of \$.705 per MHz-pop. The C licenses sold for \$0.548. The 734 A licenses sold for \$0.417.²⁴ The B licenses, on which SpectrumCo’s bids represented over 95 percent of the money spent, went for \$0.451. SpectrumCo acquired 20 MHz of spectrum covering virtually the entire country for \$2.378 billion (Figure 9). By contrast, the two largest bidders in the auction, Verizon and T-Mobile, between them acquired 40 MHz for \$6.99 billion.²⁵ Overall, the average price per MHz-pop paid by the four largest bidders in the auction apart from SpectrumCo — Verizon, T-Mobile, Cingular and MetroPCS — was \$0.68 cents. Relative to that standard, SpectrumCo saved more than \$1.1 billion.

6 Conclusion

Keynes (1936) famously concluded in the *General Theory*, “Practical men, who believe themselves to be quite exempt from any intellectual influence, are usually the slaves of some defunct economist.” But in the modern economy, particularly in fields such as portfolio theory and auction theory, the time to implementation has shrunk. Ideas developed by economists have not only played a role in designing the market games that allocate many important resources, but they also provide the insights necessary to play these games at the highest level.

In the AWS auction, SpectrumCo’s ability to alter the relative pace of price increases of the large licenses, combined with its ability to forecast final total prices, enabled it to take two calculated risks. Its “bookends” jump-bid strategy enabled it to discover that the cost of assembling a national footprint using major REAG licenses would likely become more than it was willing to pay. The strategy also forced SpectrumCo’s most direct competitor, Wireless DBS, into making billion dollar decisions with just hours of notice. SpectrumCo’s ability to forecast total auction revenue gave it the confidence that it could assemble a large number

²⁴The low prices overall for the A licenses mask significant differences between the large and small markets. For example, in the five top markets (New York, Chicago, Los Angeles, Washington, and Philadelphia) the B licenses were 21 percent cheaper than the A licenses (9 percent greater cost for areas covering 38 percent more people) but in the remainder of the country the A licenses were roughly 15 percent cheaper.

²⁵Verizon spent \$2.81 billion almost exclusively on the REAG bands, paying an average price of \$.731. T-Mobile spent \$4.18 billion, 70 percent on REAG licenses and half the rest on the two most valuable A licenses (New York and Chicago). A collection of cheaper small licenses reduced its average cost to \$.630.

of smaller licenses into a national footprint within its available budget, making bidding on these licenses a good calculated risk in spite of the exposure problem.

In the later part of the auction, SpectrumCo's strategy of steadfast bidding on its network of B licenses may have encouraged other bidders to devote most of their attention to other blocks, although tactically other large bidders could have taken actions to raise SpectrumCo's costs. The net result was a national wireless footprint at a billion dollar discount relative to competitor's prices. While opportunities to achieve such successes are hard to come by and the fine details of every auction are different, the experience suggests a value to economic and game theoretic analysis in complex auctions.

References

- Ausubel, Lawrence and Paul Milgrom, "The Lovely But Lonely Vickrey Auction," in Cramton, Shoham and Steinberg, *Combinatorial Auctions*, 2006.
- Avery, Christopher, "Strategic Jump Bidding in English Auctions," *Review of Economic Studies*, April 1998, 65(2), pp. 185-210
- Brusco, Sandro and Guiseppe Lopomo. 2002. "Collusion via Signalling in Simultaneous Ascending Bid Auctions with Heterogeneous Objects, with and without Complementarities," *Review of Economic Studies*, 69(1), pp. 1-30.
- Cramton, Peter, Yoav Shoham, and Richard Steinberg, *Combinatorial Auctions*. Cambridge, MA: MIT Press, 2006.
- Day, Bob and Paul Milgrom, "Core-Selecting Auctions," *International Journal of Game Theory*, 36, 2008, 393-407.
- Gul, Faruk and Ennio Stacchetti, "The English Auction with Differentiated Commodities," *Journal of Economic Theory*, 2001.
- Hegeman, John, "A Dominant Strategy Mechanism for Bidders with Budget Constraints," Stanford Working Paper, 2008.
- Keynes, John Maynard, *The General Theory of Employment, Interest and Money*, New York: Harcourt, Brace and Company, 1936.
- Klemperer, Paul, "What Really Matters in Auction Design," *Journal of Economic Perspectives*, 2002.
- Milgrom, Paul, "Putting Auction Theory to Work: The Simultaneous Ascending Auction." *Journal of Political Economy*, 2000, 108(2), pp. 245-72.
- Milgrom, Paul, *Putting Auction Theory to Work*, Cambridge: Cambridge University Press, 2004.
- Milgrom, Paul, "Package Auctions and Package Exchanges," *Econometrica*, 2007, 75(4), pp. 935-966.
- Rothkopf, Michael, "Thirteen Reasons Why the Vickrey-Clarke-Groves Process is Not Practical," *Operations Research*, 2007: 55(2): 191-197.
- Zheng, Charles. "Jump Bidding and Overconcentration in Decentralized Simultaneous Ascending Auctions," Iowa State Working Paper, 2006.

Table 1: Prices Paid by the Five largest Buyers in AWS Auction

Bidder	Total Winning Bids	MHz-Pops	Price per MHz-Pop
SpectrumCo	\$2,377,609,000	5,267,189,470	\$0.45
Cingular	1,334,610,000	2,436,458,880	\$0.55
T-Mobile	4,182,312,000	6,638,718,070	\$0.63
Verizon	2,808,599,000	3,840,952,220	\$0.73
MetroPCS	1,391,410,000	1,445,444,020	\$0.96
Four incumbents	9,716,931,000	14,361,573,190	\$0.68

Notes: (1) A MHz-pop is a standard unit for spectrum, with one unit meaning one MHz of bandwidth covering one person. (2) The high prices paid by MetroPCS reflect to some extent the fact that their licenses covered dense urban areas; these licenses typically sell for a premium in per-unit term. (3) SpectrumCo was an entrant.

Table 2: Spectrum Price per MHz-Pop in Past Auctions

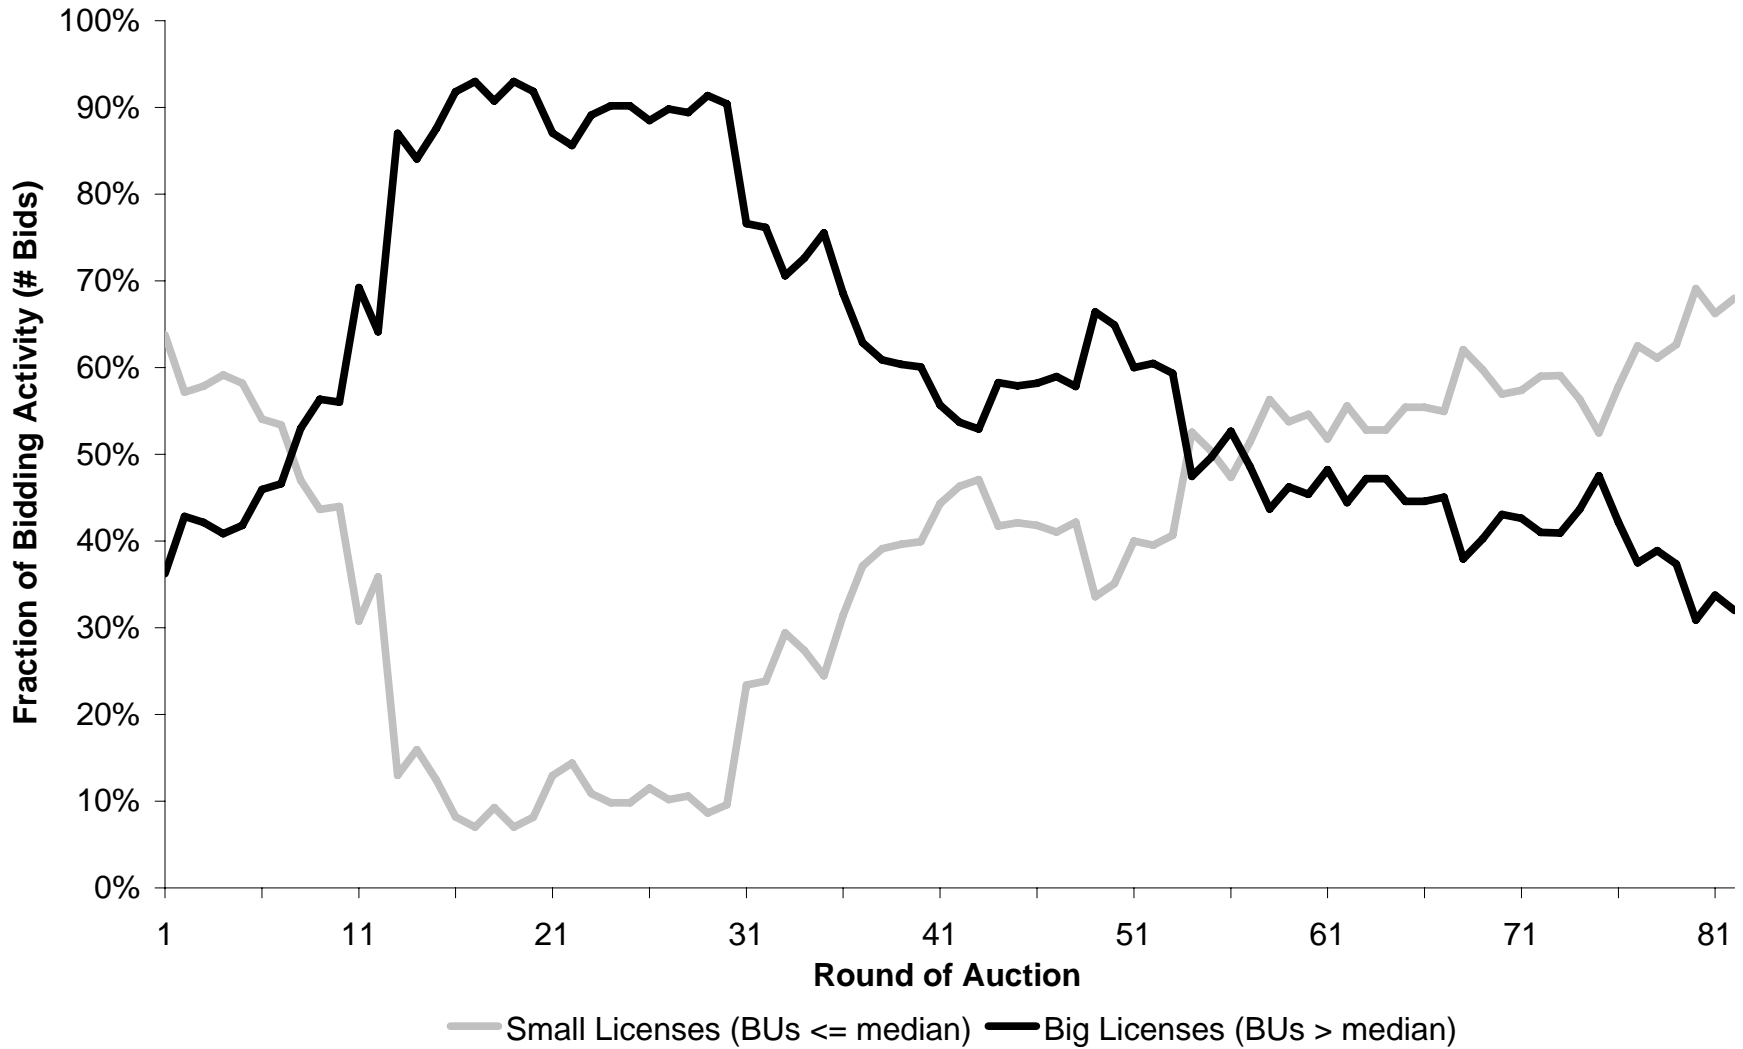
Auction	Description	Date	Revenue (US \$m)	Price per MHz-Pop
5	PCS C Block	May 1996	13,429	1.77
10	PCS C Block Re-Auction	July 1996	697	1.50
11	PCS D,E,F Block	Jan 1997	2,716	0.36
22	PCS	March 1999	533	0.20
34	800 MHz Auction	Sept 2000	337	0.18
35	PCS C&F Block	Jan 2001	17,596	4.37
58	Broadband PCS	Jan 2005	2,250	1.05
66	AWS Auction	Sept. 2006	13,879	0.54
73	700 MHz	March 2008	19,592	1.11

Table 3: Exposure and Revenue Rise in Major Auctions

FCC Auction	Total Rounds in Auction	First round with Exp. ≥ 90% of Final Revenue	PWB/Final Revenue in First Round with Exp. ≥ 90% of Final Revenue
22	78	6	0.49
30	73	14	0.52
33	66	11	0.16
34	76	20	0.41
35	101	12	0.21
37	62	6	0.16
44	84	21	0.82
53	49	11	0.33
58	91	4	0.38
66	161	10	0.25

Note: PWB stands for "Provisionally Winning Bids" or current revenue in a given round.

Figure 1: Pattern of Bidding Activity in Auction 66



Note: Plot ends at round 82, when there were 75 bids. Beyond this point, there were less than 50 bids per round and the plotted fractions fluctuate

Figure 2a: Round of Final Bid by License Size (Auction 35)

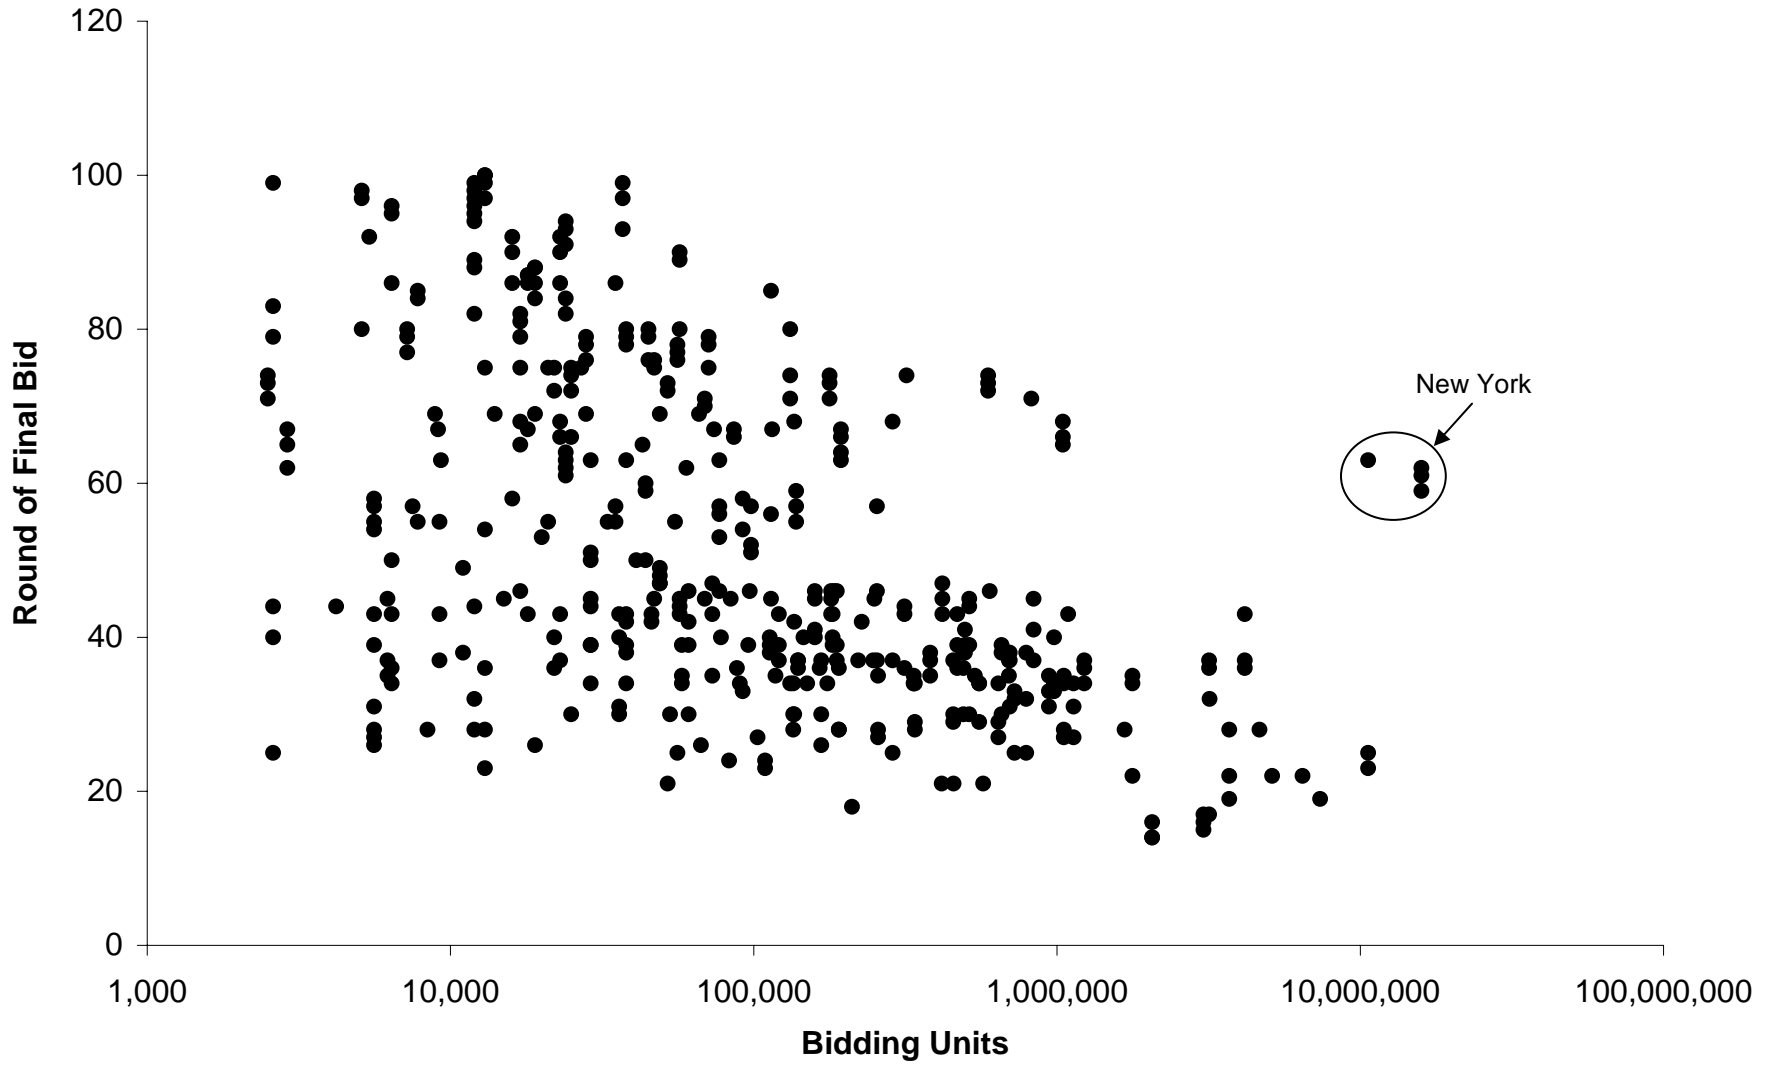
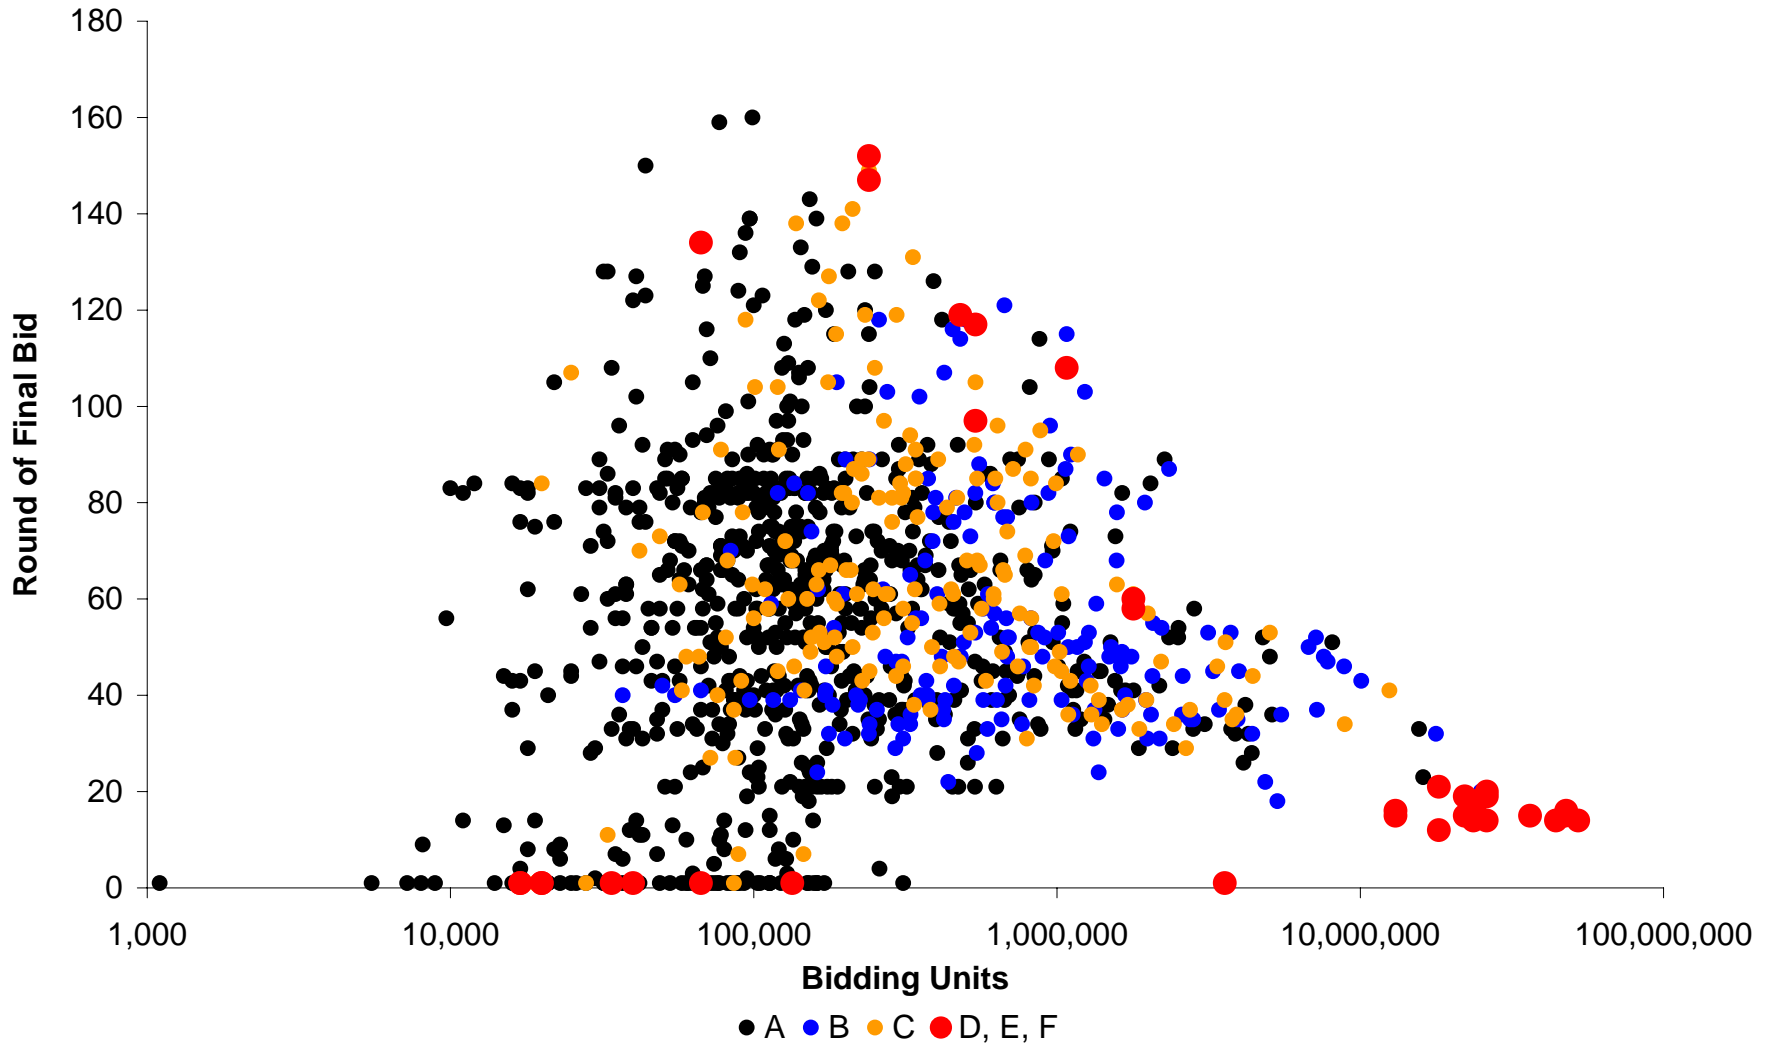


Figure 2b: Round of Final Bid by License Size (Auction 66)



Note: Blocks D,E,F were divided into 12 REAG licenses; blocks B and C were divided into 176 EA licenses; block A into 734 CMA licenses.

Figure 3: Revenue and Exposure in Auction 35

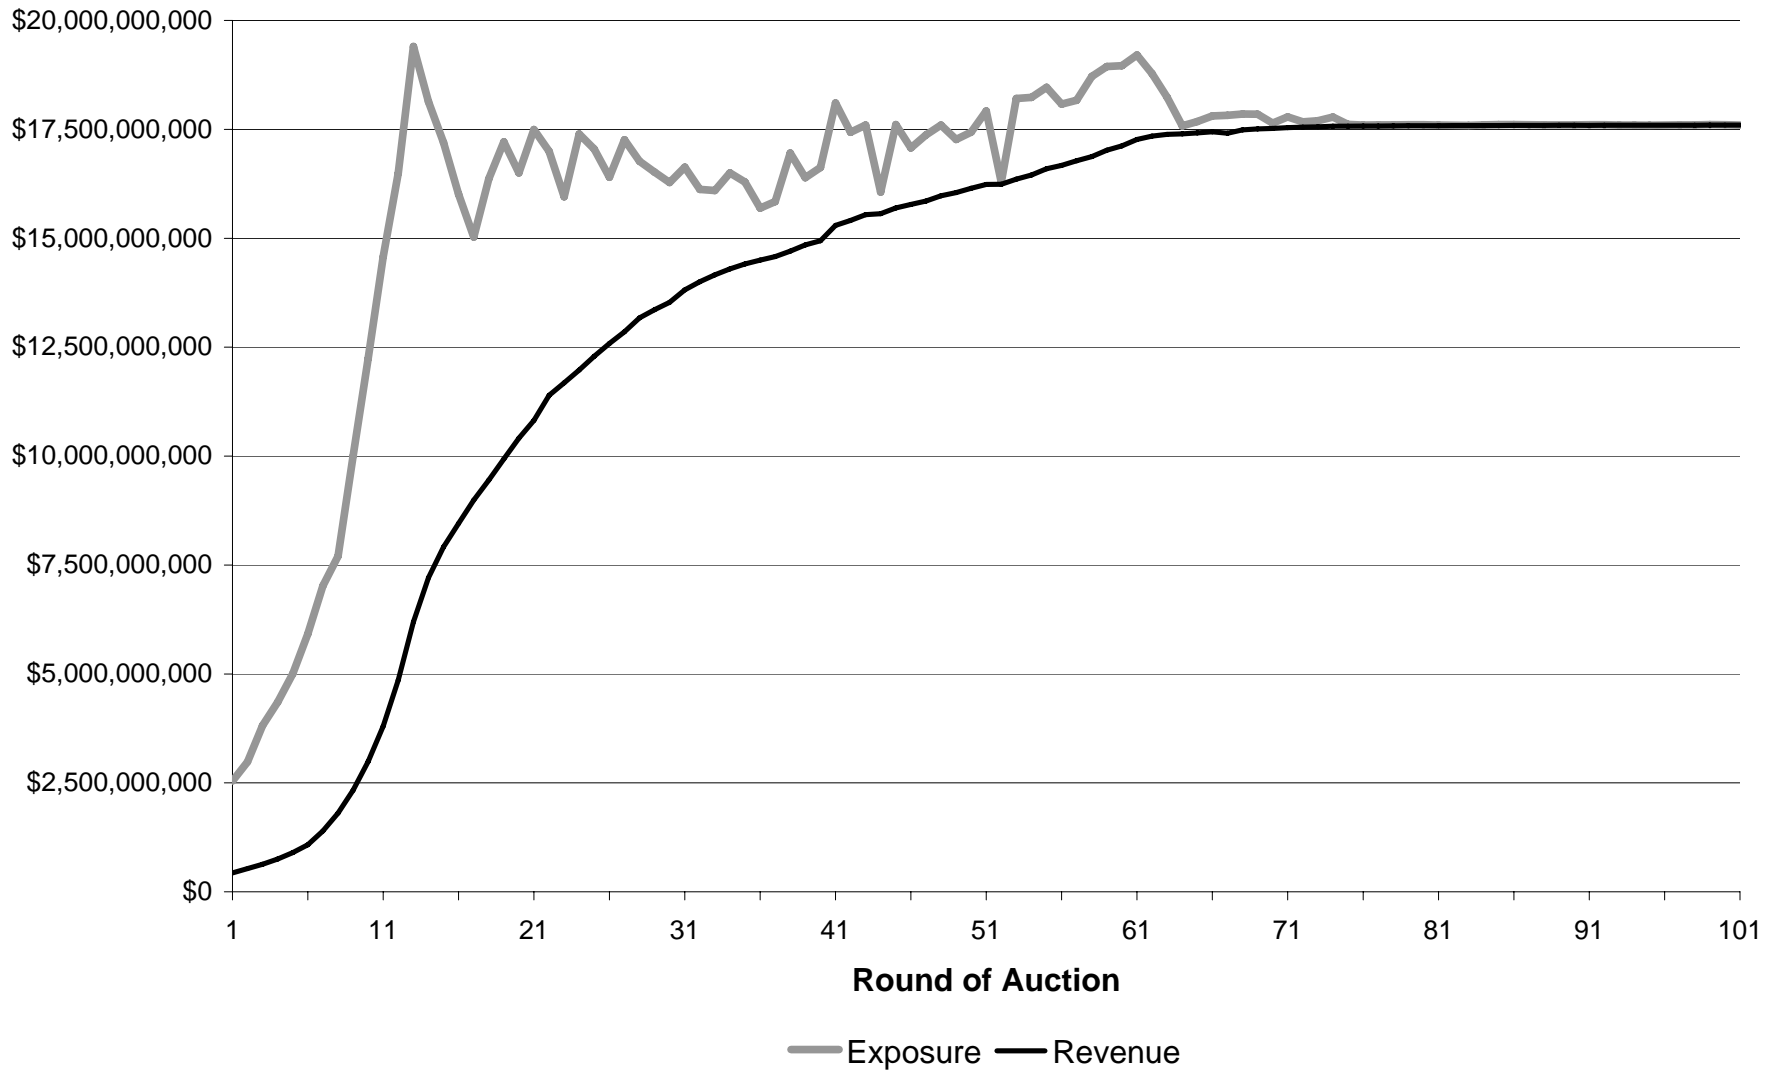


Figure 4: Revenue and Exposure in Auction 66

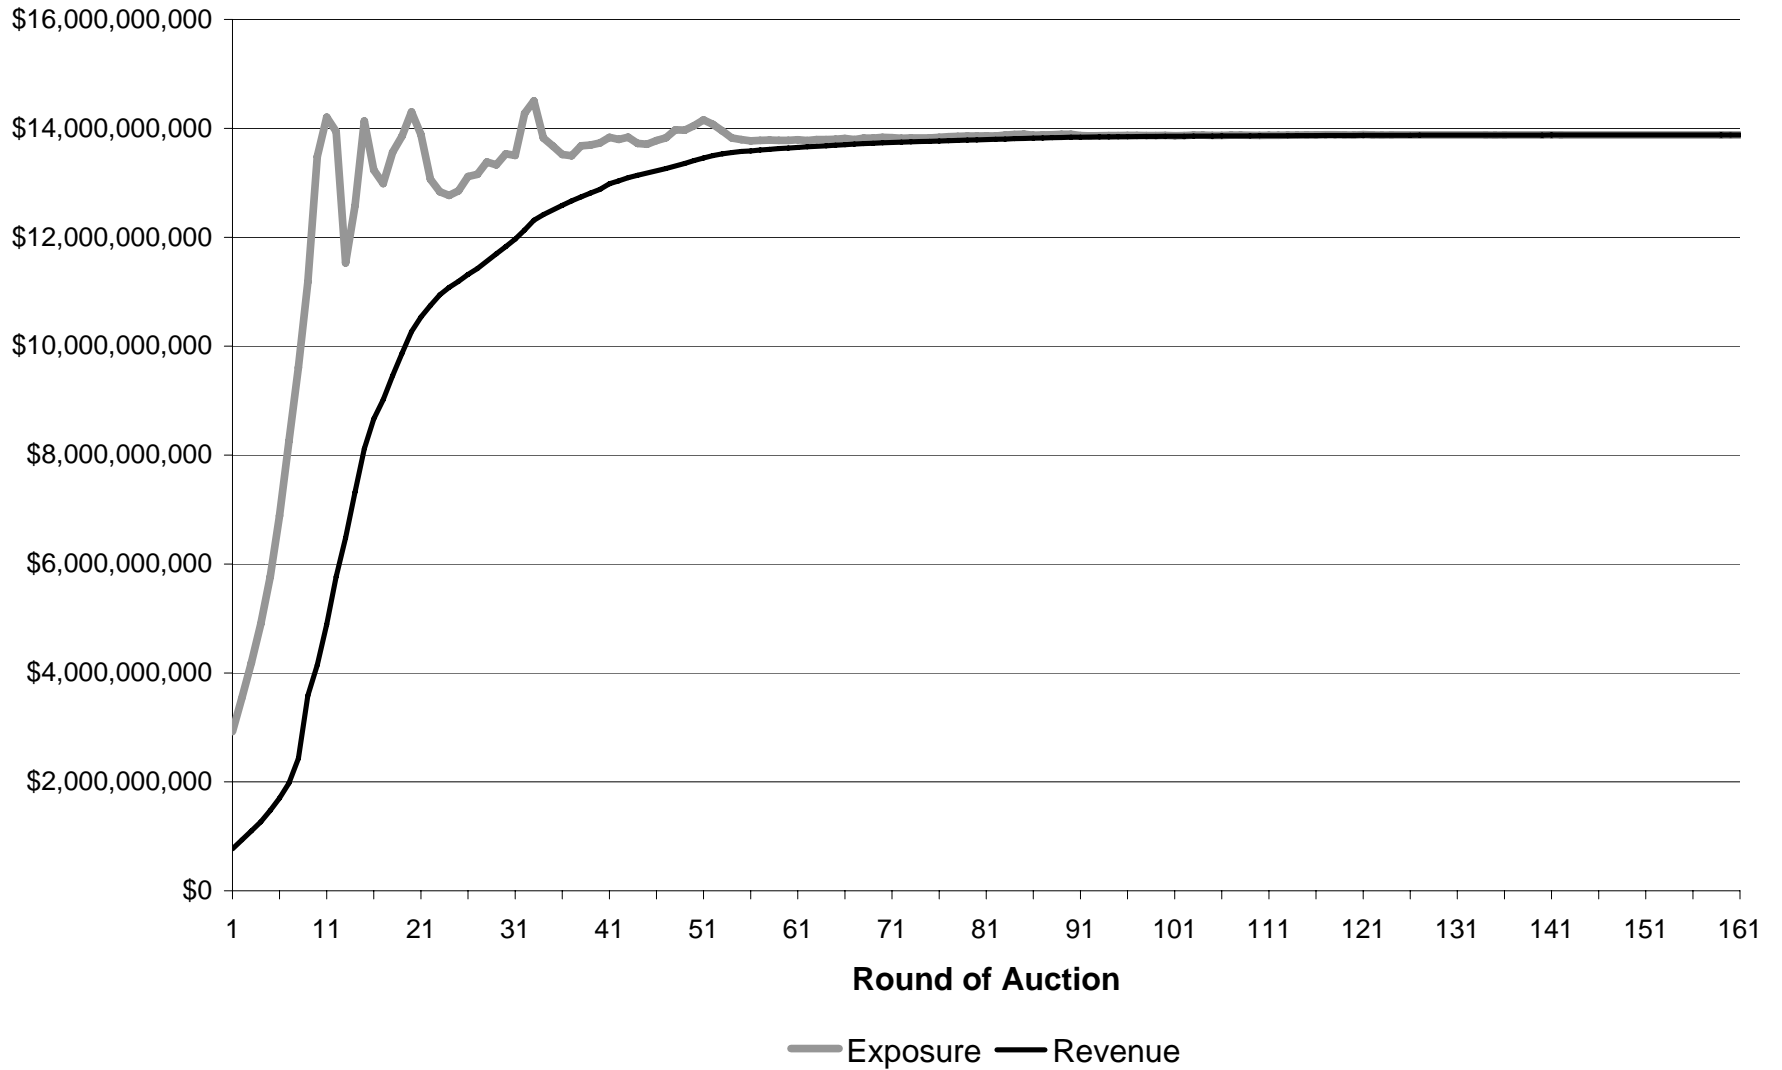


Figure 5a: Bidder Exposure in Auction 35

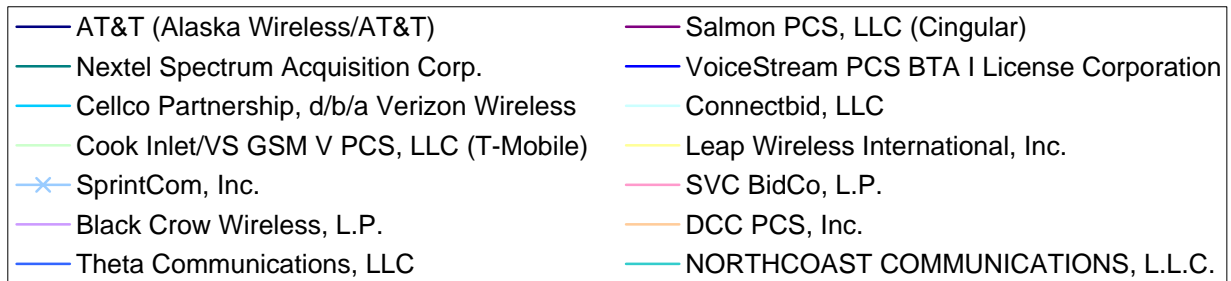
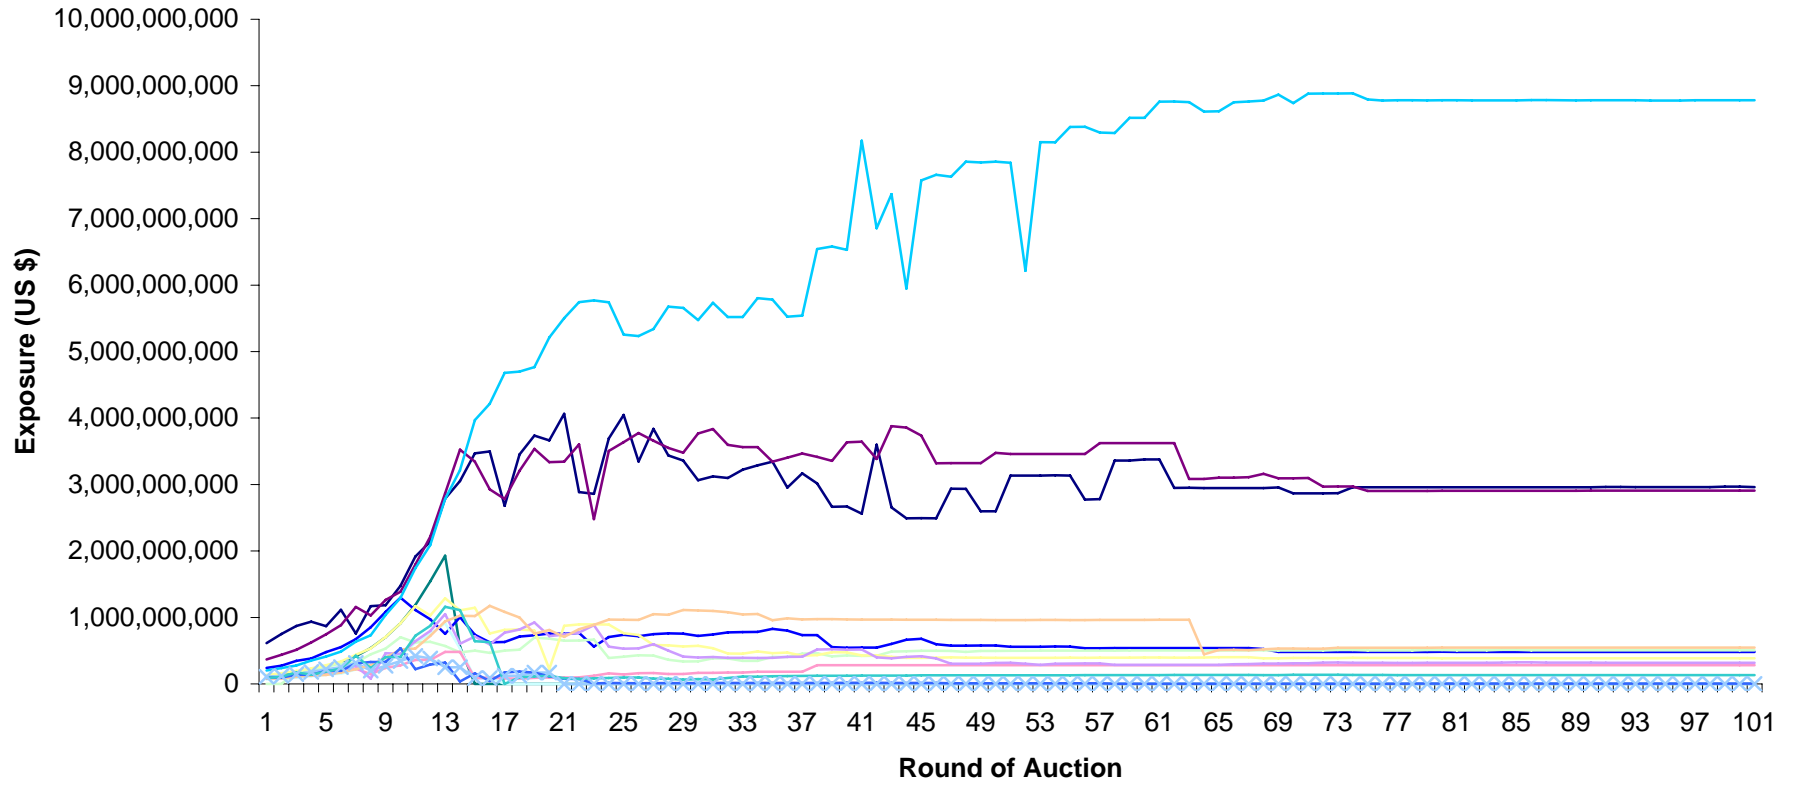


Figure 5b: Bidder Exposure in Auction 66

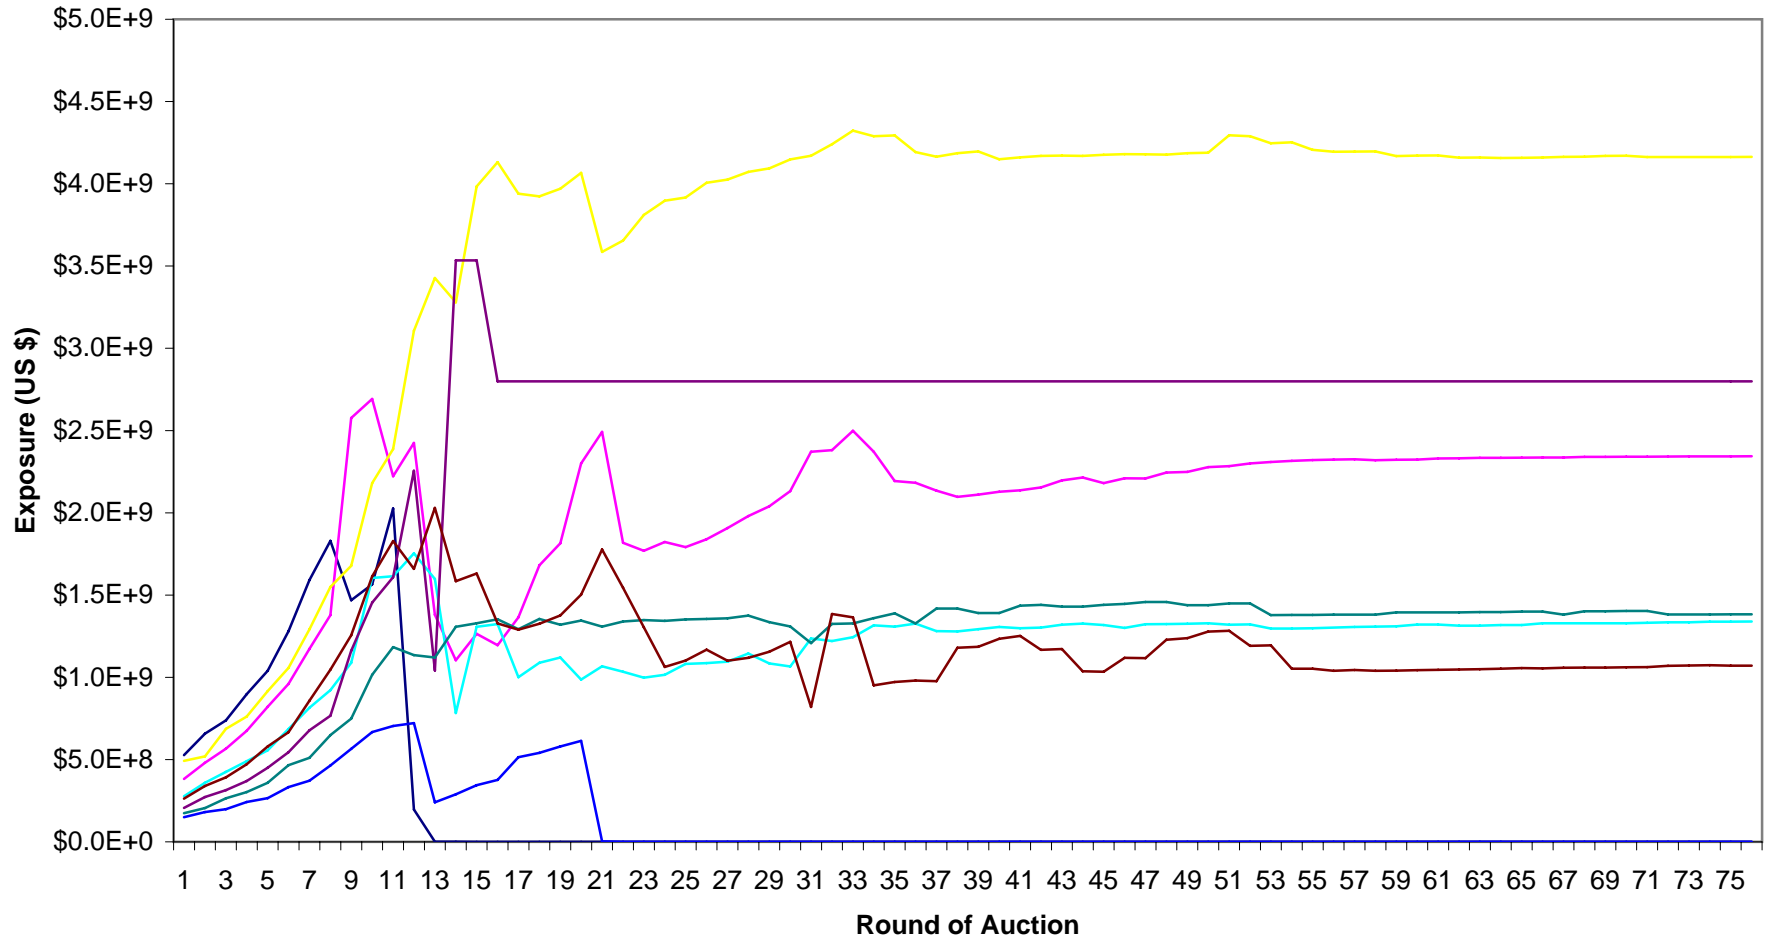


Figure 6: Ratio of Maximum Exposure to Final Gross Revenue in Major FCC Auctions

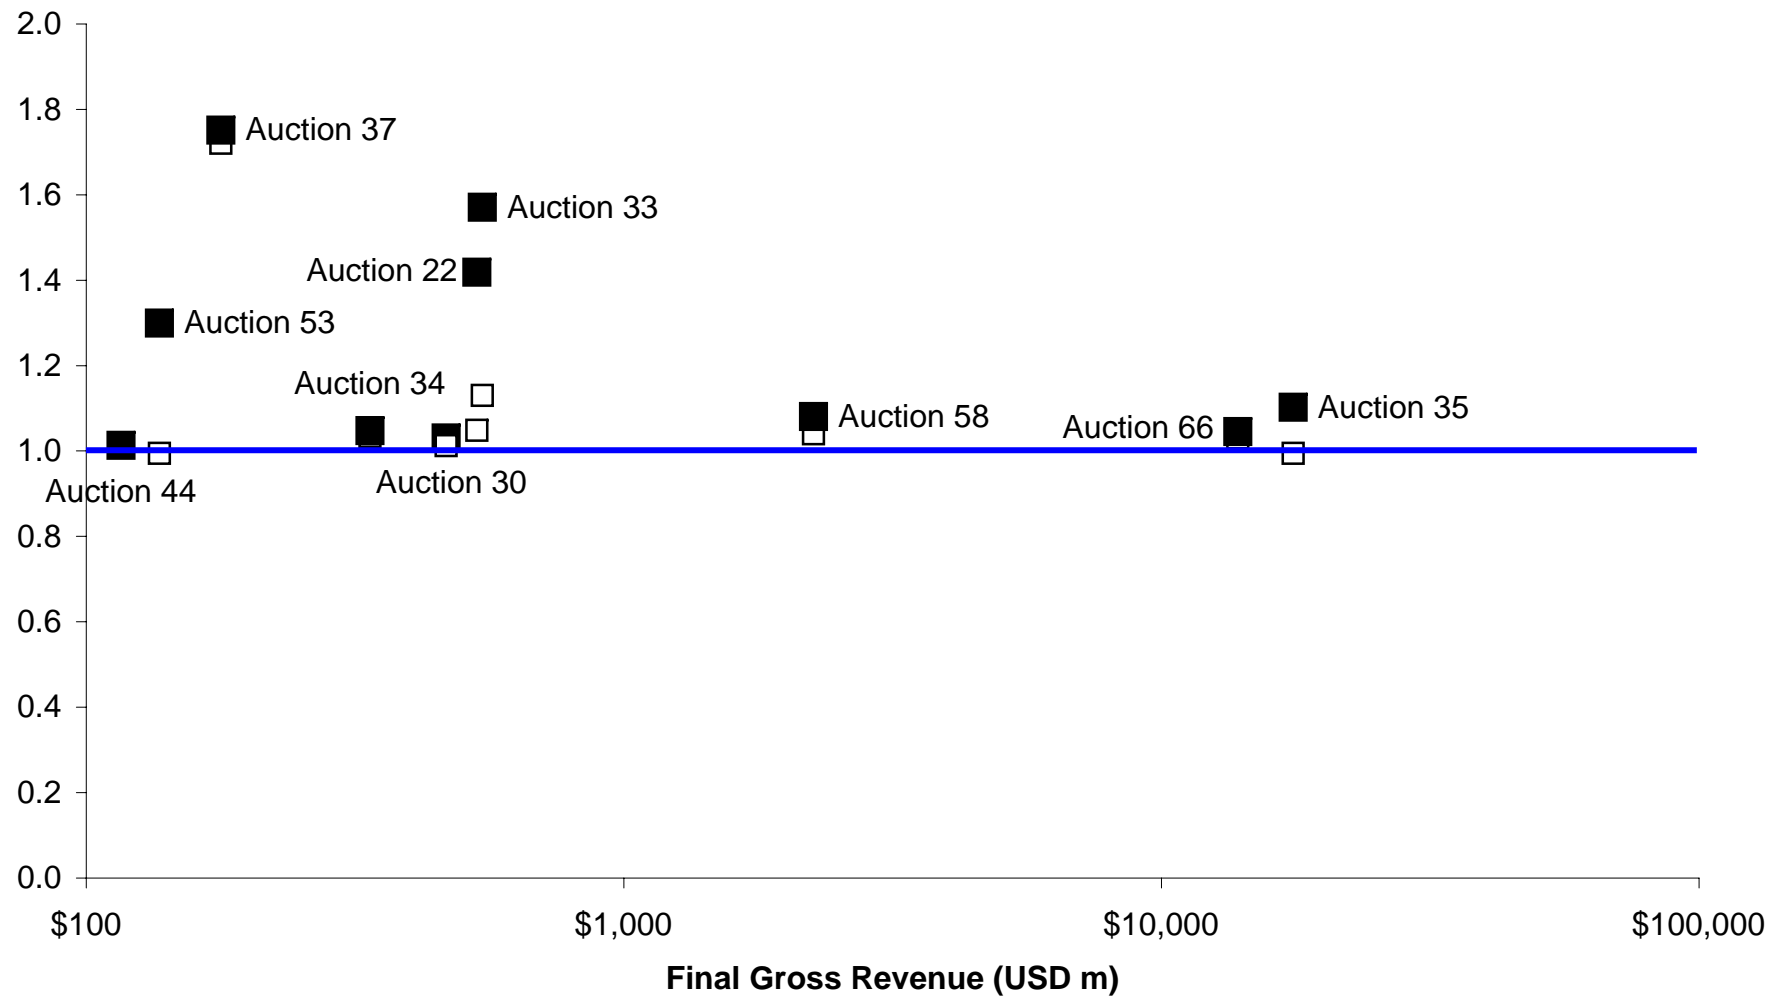


Figure 7: Optimal Price Management

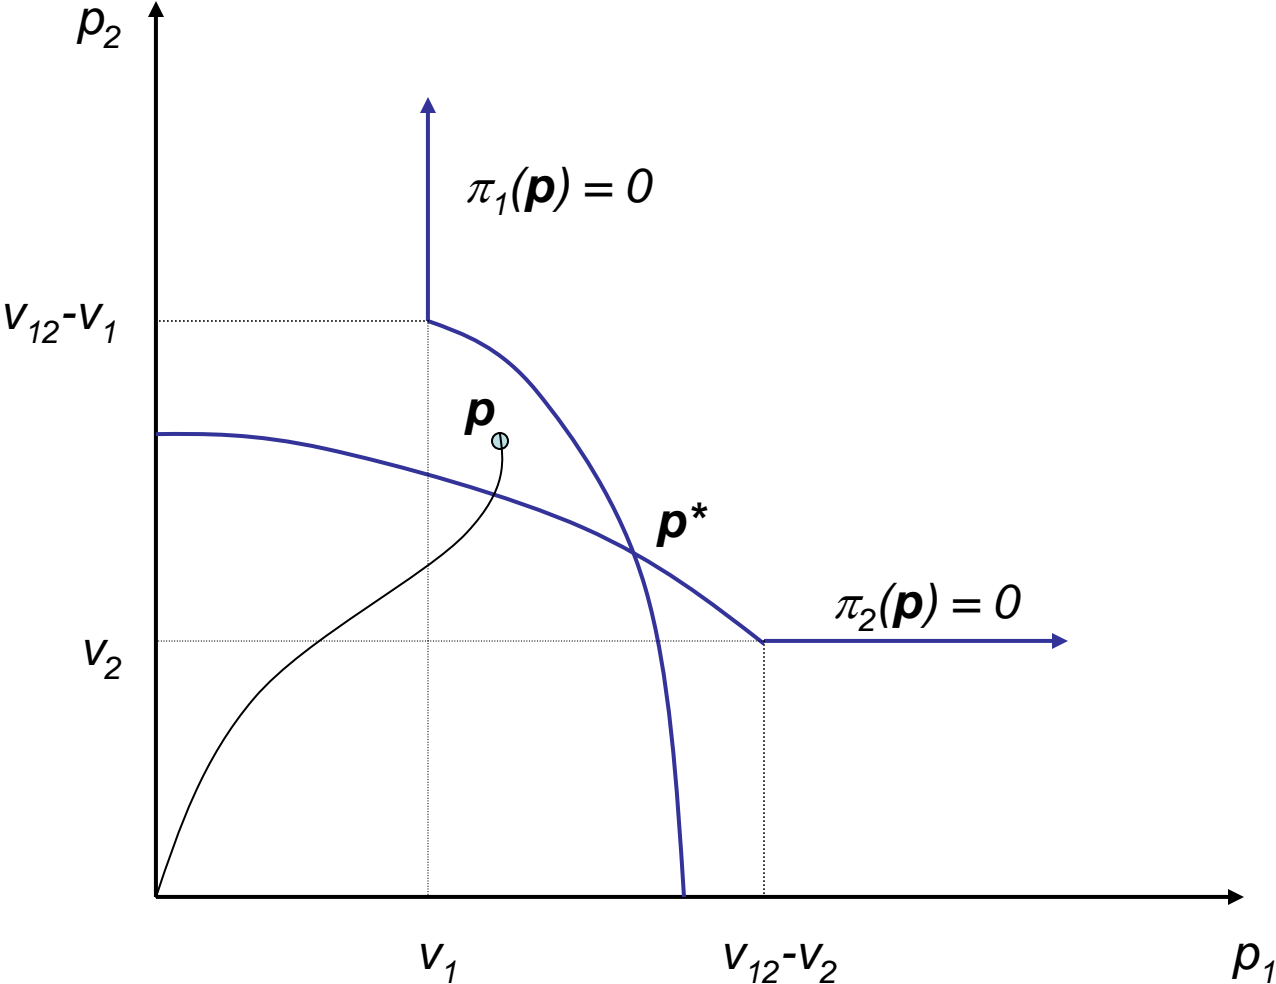


Figure 8: Price across AWS Bands

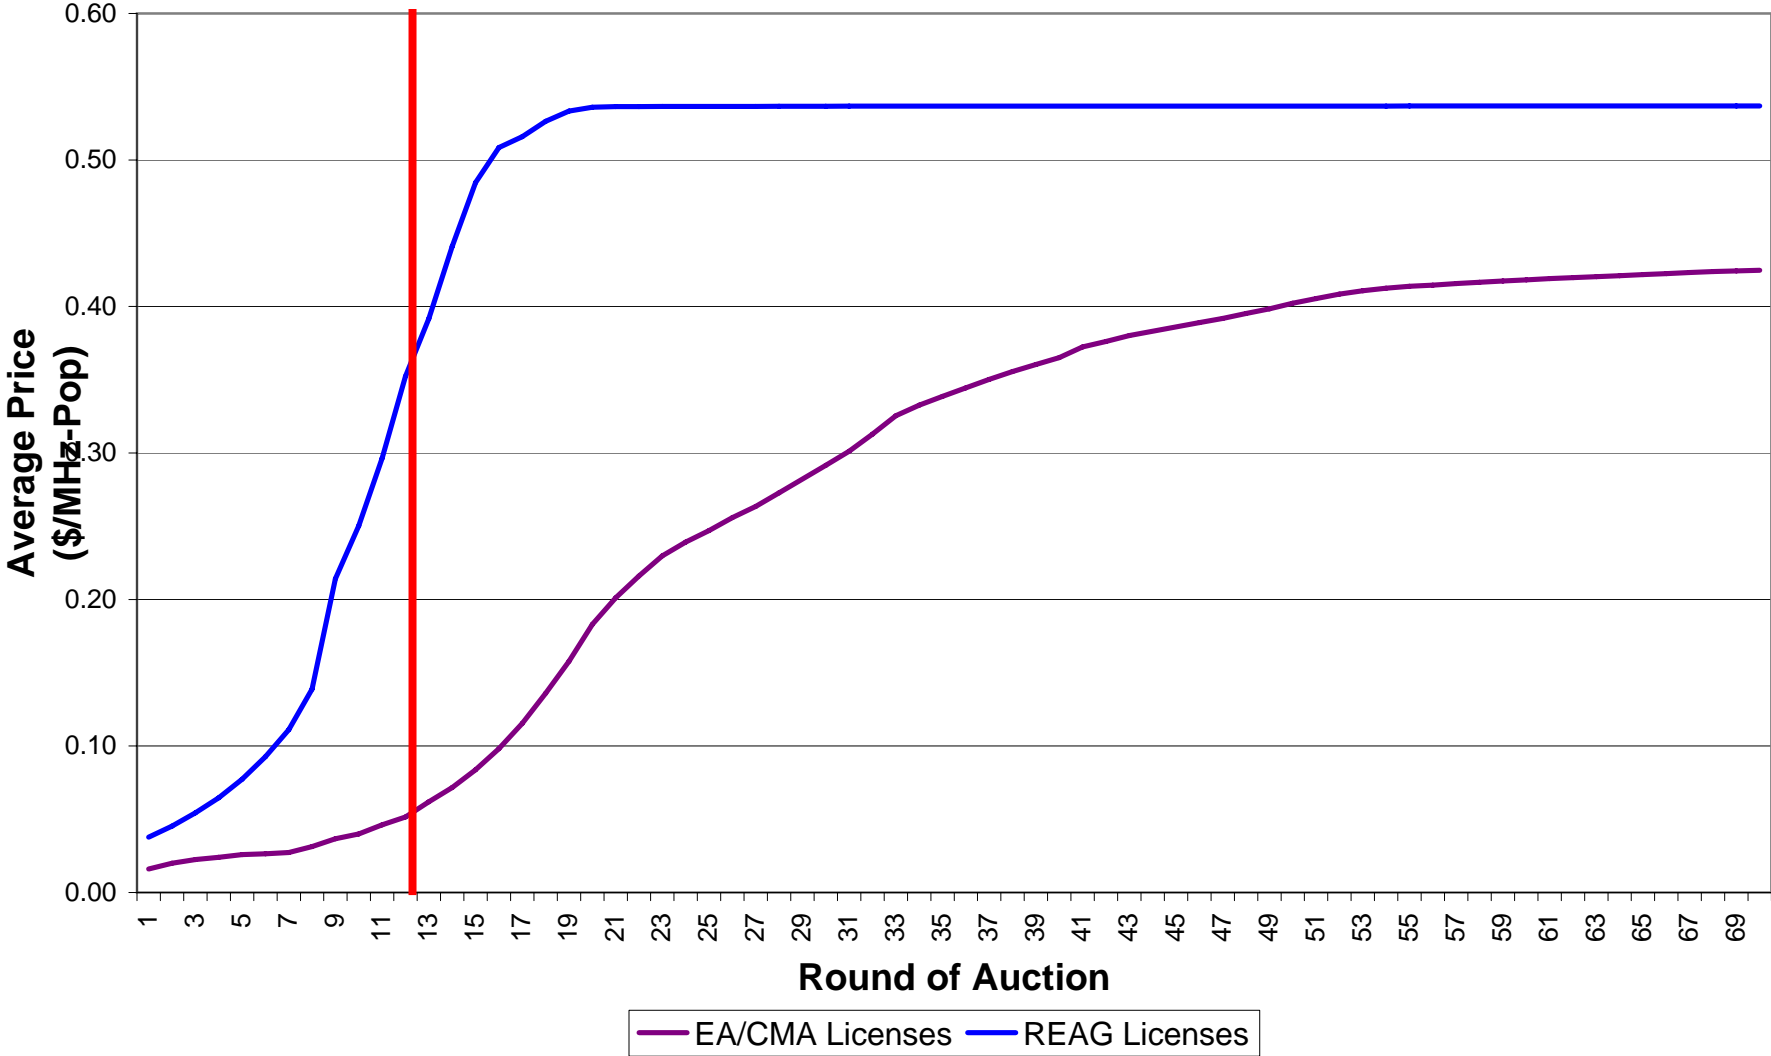


Figure 9: SpectrumCo's Purchases in AWS Spectrum Auction

