Market Design and the FCC Incentive Auction*

Lawrence M. Ausubel†, Christina Aperjis‡ and Oleg Baranov§

Current Draft: 15 July 2017

Preliminary and Incomplete

Not for Quotation

Abstract

Market-based mechanisms for the allocation of spectrum licenses have been a prominent feature of the telecommunications landscape for more than two decades, but these mechanisms have taken as given that the offered spectrum will be cleared of prior users. In March 2017, the FCC completed the Incentive Auction, the first-ever auction that incorporated voluntary clearing directly into the allocation mechanism. This article seeks to examine the results of the FCC Incentive Auction from a market design perspective.

---

* *Disclaimer:* This disclaimer is required by Contract No. FCC14C0010. The Federal Communications Commission (FCC) has not reviewed or approved any statement in this document for accuracy or validity. The FCC and its employees do not endorse goods or services provided by Power Auctions LLC or any other firm, except as allowed by 5 C.F.R. 2635.702(c)(1)-(2), which do not apply here.

† Department of Economics, University of Maryland, College Park, MD 20742, USA. Email: ausubel@econ.umd.edu

‡ Power Auctions LLC, 3333 K St NW Suite 425, Washington, DC 20007, USA. Email: aperjis@powerauctions.com

§ Department of Economics, University of Colorado, Boulder, CO 80309, USA. Email: oleg.baranov@colorado.edu

1. Introduction

Market-based mechanisms for the allocation of spectrum licenses have been a prominent feature of the telecommunications landscape for more than two decades. Federal Communications Commission (FCC) Auction #1 (the Nationwide Narrowband Auction) of July 1994 and FCC Auction #4 (the Broadband PCS A and B Block Auction) of December 1994 – March 1995 were pioneering United States efforts, while UMTS/IMT-2000 Auctions in the UK (March – April 2000) and Germany (July – August 2000) are best known for obtaining stunningly high prices for 3G licenses.[1]

However, spectrum is never "new" and, typically, some incumbent party has been using a given swath of spectrum before it is auctioned as a spectrum license. Current practice is generally to make termination of the prior use of the spectrum a mandatory part of the allocation process, through the issuance of regulations (and sometimes with payment of some amount of compensation to the incumbent user). In recent years, economists have sought to integrate the clearing of incumbent uses of the spectrum into the auction process. Integrating voluntary clearing into spectrum auctions offers at least two potential advantages over current practice. First, to the extent that the cleared incumbents are the "winners" of an auction process, they may be incentivized to terminate their prior use in an expeditious way, instead of filing lawsuits or otherwise delaying the clearing. Second, by utilizing such a process, the policymaker can be confident that the repurposing of spectrum among uses and users is creating social surplus—and indeed the amount of repurposing can be made an endogenous part of the process.

Bidding in the first-ever auction that incorporated voluntary clearing began on March 29, 2016 and concluded on March 30, 2017. Dubbed the FCC "Incentive Auction", its objective was to repurpose spectrum from television broadcasting to mobile broadband. The auction was successful in its objective, repurposing 84 MHz of low-band spectrum, including 70 MHz of licensed spectrum and 14 MHz for unlicensed use.

This article seeks to examine the results of the Incentive Auction from a market design perspective. In turn, we will review and assess the effectiveness of the meta-design of the auction, followed by looking at particular facets of the reverse auction and the forward auction.

At the outset, it should be observed that spectrum auctions are often put forward as a "poster child" for market design. While allocating spectrum via auction undoubtedly benefits government revenues, there has always been ambiguity as to whether previous auctions were successful in creating and maximizing social welfare. However, the FCC Incentive Auction presents a stronger case: the new users revealed their values to be at least $19.3 billion, while the prior users revealed their costs to be at most $10 billion. In our view, it provides a compelling case study of market design improving social welfare.

The account that we provide is most certainly an insider account. Larry Ausubel was a member of the design team and an author of the key design documents.[2] Christina Aperjis participated in many of the

[1] Accounts and analyses of some of the early spectrum auctions are provided in McMillan (1994), Cramton (1995), McAfee and McMillan (1996), Jehiel and Moldovanu (2001) and Binmore and Klemperer (2002).

[2] Milgrom, P., L. Ausubel, J. Levin and I. Segal (2012), "Incentive Auction Rules Option and Discussion," Appendix C to *Expanding the Economic and Innovation Opportunities of Spectrum Through Incentive Auctions*, Notice of Proposed Rulemaking, Docket No. 12-268, 27 FCC Rcd 12357 (2012), https://apps.fcc.gov/edocs_public/attachmatch/FCC-12-

design decisions, was an author of many of the technical appendices to the FCC's public notices, and spent too much of 2016–2017 in the FCC auction war room. Oleg Baranov also participated in the design of the assignment phase. And Power Auctions LLC, our company affiliation, was hired by the FCC to do the software implementation of the Incentive Auction.

The article is structured as follows. Section 2 provides an overview of the Incentive Auction. Section 3 looks at the efficiency of the overall design and particular facets of the reverse auction. Section 4 examines the main phase of the forward auction. Section 5 identifies some interesting features of the assignment phase. Finally, Section 6 concludes.

2. Overview

The Incentive Auction comprised two separate but interdependent auctions: a reverse auction and a forward auction. The reverse auction was a descending clock auction in which television broadcasters bid to relinquish their spectrum usage rights. The forward auction consisted of an ascending clock auction phase for generic blocks, followed by an assignment phase to map generic blocks to physical frequencies.[3]

Each *stage* of the auction was associated with a clearing target. In each stage, the reverse auction was conducted first to determine the cost (in terms of payments to broadcasters) to achieve the clearing target of the stage. Then, the forward auction was conducted to determine whether the current stage would be the final stage of the auction. A stage was deemed to be final when the final stage rule was met. The final stage rule requires that forward prices are competitive and that net forward auction proceeds are sufficient to cover the incentive auction costs. The auction concluded after 4 stages.

The biggest surprise in the incentive auction was that only one of the four nationwide wireless providers was bidding in the auction until the conclusion of the clock phase. Sprint announced before the auction that it would not participate. Verizon applied to participate in the auction and became a qualified bidder but did not bid in the auction and, as a result, lost its eligibility to bid after the initial round of the auction. AT&T was bidding in the first rounds of the auction, but significantly reduced its requested commitment in the last few rounds of stage 1 and bid to drop out of the auction completely throughout stages 3 and 4.[4] T-Mobile was the only nationwide wireless provider with significant winnings. Other bidders with significant winnings in the auction were Dish and Comcast.

[More to be added]

118A2.pdf, and Milgrom, P., L. Ausubel, J. Levin and I. Segal (2013), "Auctionomics/Power Auctions Option for Forward Auction," Supplement to Appendix C, https://ecfsapi.fcc.gov/file/7022116356.pdf.

[3] Only about 0.7% of Incentive Auction proceeds came from the assignment phase, suggestive that all seven blocks were viewed by bidders as nearly equivalent and that it was appropriate to treat all blocks as generic spectrum in the clock phase.

[4] The auction rules did not allow AT&T to drop out of the auction completely. AT&T's commitment at the conclusion of the clock phase was $894,137,000. Even though AT&T was not able to drop out of the auction completely, its clock phase commitment was significantly lower than its maximum commitment during stage 1 which was $7,369,594,000 (after round 21).

3. The Overall Design of the Incentive Auction

3.1 Efficient Clearing

The incentive auction mechanism allowed the quantity of cleared, licensed spectrum to be any number up to ten (paired) blocks. Based on all publicly-available information, the auction achieved exactly the efficient amount of clearing.

The cost of clearing spectrum for the seven-block band plan (84 MHz) was $10.05 billion whereas the cost of clearing spectrum for the eight-block band plan (108 MHz) was revealed to be $40.31 billion (see Table 8). Thus, the incremental cost of clearing an eighth block would have been $30 billion, which is more than the buyers' combined bids to purchase the first 7 blocks. In the terminology of spectrum valuation, the cost of clearing an eighth block would have been almost $10 per MHz per population ($10 per MHz-pop), which is extremely high.

Meanwhile, the resulting seven blocks are unimpaired, contiguous to one another, and configured well for international alignment with Canada and Mexico.

3.2 Reverse Auction Design

The reverse auction used a descending clock format. In each round, each station was given a price offer for one or more relinquishment options. The possible relinquishment options were to go off-air or move to a different band; for instance, a UHF station could bid to move to the High-VHF band or to the Low-VHF band. Price offers for a given station descended from one round to the next. Given the price offers for the round, a bidder could then select its preferred relinquishment option for its station or bid to drop out of the auction. If, after a round, the auction system could not find a feasible assignment for a station in its pre-auction band subject to various interference constraints, the station was deemed provisionally winning. The winners of the reverse auction were the stations that became provisionally winning in the final stage of the auction.

The reverse auction appears to have performed well in terms of inducing bidders to participate in the auction. The FCC has not published the names or the number of stations that made initial commitments to participate in the auction, but the fact that the initial clearing target was 126 MHz suggests that the number of participating stations was large.

The reverse auction also performed well in terms of clearing 84 MHz at a reasonable cost. However, the cost for clearing more than 84 MHz appears to be very large. Specifically, the incremental cost in going from 84 MHz to 108 MHz (8^{th} block) was $30 billion, the incremental cost in going from 108 MHz to 114 MHz (9^{th} block) was $14 billion, and the incremental cost in going from 114 MHz to 126 MHz (10^{th} block) was $32 billion (see Table 8).

The cost of clearing more than 7 blocks is partly explained by increasing marginal costs and the fact that the larger clearing targets had to accommodate Channel 37. Channel 37 is a frequency region set aside for Radio Astronomy that could not be repacked during the incentive auction. As shown in Figure 1 below, clearing targets above 84 MHz had to clear additional spectrum compared to the 84 MHz clearing targets because of the required guard bands around Channel 37 and between the TV band and the wireless band.

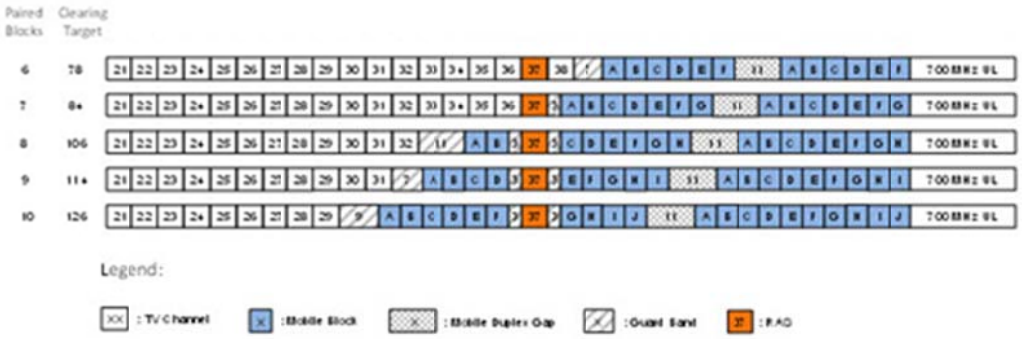


Figure 1: 600 MHz Band Plans

Strategic bidding by reverse auction bidders could have also contributed to increasing the clearing costs. The effect of strategic bidding in the clearing cost can be larger when the FCC attempted to clear more spectrum, as it did in the first three stages of the auction.

One of the main objectives of the reverse auction design was obvious strategy-proofness (Milgrom & Segal, 2015). Since it is not possible to achieve obvious strategy-proofness for all types of bidders, the goal was reduced to achieving obvious strategy-proofness for single-minded bidders. In particular, the design is obviously strategy-proof for bidders that owned a single station and that only selected one relinquishment option (e.g. to go off-air) when applying to participate in the auction. However, the design is not strategy-proof for a bidder that owns multiple stations nor for a bidder that owns a single station but selected multiple relinquishment options on its application.

We next provide evidence suggesting that some bidders with multiple stations might have strategically reduced supply in the auction. We then discuss why a bidder that can bid on multiple relinquishment options might find beneficial to deviate from truthful bidding. We conclude that the reverse auction design, which focused on obtaining an obviously strategy-proof mechanism for single-minded bidders, could have paid greater attention to limiting the gaming by multi-minded bidders.

Strategic Supply Reduction

A bidder that owns multiple stations in a certain market might find it beneficial to strategically reduce its supply in the auction. For example, consider a bidder that owns two stations that interfere with each other. By dropping out of the auction for one station early on (or by only signing up one station for the auction), the bidder may be able to cause its other station to become provisionally winning at a high price. Strategic supply reduction can significantly increase the clearing cost (Doraszelski, et al., 2016).

We note that strategic supply reduction is not risk-free as the bidder does not know which other stations are bidding in the auction and thus cannot be sure that dropping one of its stations out of the auction will cause its other station to become provisionally winning. Still, it is a strategy that bidders could employ, along the lines of a demand reduction strategy in a forward auction (Ausubel, et al., 2014). Some evidence of potential supply reduction in the reverse auction is provided by the auction results for spectrum speculators that acquired TV licenses in anticipation of the incentive auction.

OTA Broadcasting is a private equity firm founded by Michael Dell's MSD Capital that acquired a number of TV licenses before the auction. According to its website (http://otabroadcasting.com/), OTA Broadcasting owned 24 stations before the auction. 10 of these stations were sold in the incentive auction. It is plausible that OTA Broadcasting employed supply reduction strategies in Pittsburgh, San Francisco, and Charlotte. In Pittsburgh, OTA Broadcasting owned 11 stations before the auction but sold only 5 of those stations. In San Francisco, OTA Broadcasting owned 2 stations before the auction and sold 1 of those stations. In Charlotte, OTA Broadcasting owned 2 stations before the auction and sold 1 of those stations in the incentive auction.

Finally, OTA Broadcasting owns a station in Wilkes Barre-Scranton, PA and also owned a station in New York before the auction. Though these two stations do not interfere with each other, the Wilkes Barre-Scranton station interferes with several stations in New York. OTA Broadcasting did not sell its Wilkes-Barre station but did not sell its New York station in the auction. It is plausible that by not selling the Wilkes-Barre station, OTA attempted to cause the UHF band to be more congested in New York, which in turn caused its New York station to become provisionally winning when other New York stations dropped out of the auction, earlier than it would have otherwise.

NRJ TV LLC, another private equity firm that acquired TV licenses before the incentive auction, potentially conducted supply reduction in Los Angeles and San Francisco. In each of those markets, NRJ TV owned two or more stations but only sold one station in the auction.

Strategic Bidding with Multiple Relinquishment Options

A second-order effect with respect to strategic behavior is the absence of strategy-proofness for a single-station owner that selected multiple relinquishment options on its application. A bidder could select up to three relinquishment options for each of its stations when it applied to participate in the reverse auction: Go off-air, Move to Low-VHF, and Move to High-VHF. The algorithm that calculated the price offers for each relinquishment option of a station used vacancy information. A station's vacancy in each band depended, among other things, on the station's currently held option. Thus, a bidder interested in moving to High-VHF (from UHF) could potentially achieve higher a price offer for moving to High-VHF by bidding to go off-air in the early rounds of the auction and then bidding to move to High-VHF a few rounds before it expected its station to become provisionally winning. Of course such a strategy is risky because if, in this example, the station became provisionally winning while it was bidding to go off-air in the final stage of the auction, it would have committed to go off-air and thus would be required to stop broadcasting after the auction.

3.3 Forward Auction Proceeds

In 2012, Congress authorized the FCC to conduct the incentive auction, to repurpose the 600 MHz band, and Auction 97 for Advanced Wireless Services (AWS-3). $7 billion from the proceeds from those auctions were to be used to fund the First Responded Network Authority (FirstNet) that will provide emergency responders with the first nationwide, high-speed, broadband network dedicated to public safety.

AWS-3, which was conducted in 2015-2016, generated net proceeds (net of bidding credits) of $41.33 billion, which was enough to fully fund FirstNet. As a result, when the incentive auction started it was known that it would not need to raise any funding for FirstNet. The incentive auction, which was conducted in 2016-2017, generated net proceeds of $19.31 billion. AWS-3 and the incentive auction

achieved the highest revenue and the second highest revenue, respectively, across all FCC auctions. Thus, the package of AWS-3 and the incentive auction was extremely successful, as it provided a large amount of spectrum for wireless use, repurposed the 600 MHz band, funded FirstNet, and generated a significant amount to be applied to deficit reduction.

Still, the revenues of the incentive auction have been compared (unfavorably) to the revenues of the AWS-3 auction. The significant difference in revenue between the two auctions is largely due to the different levels of participation by AT&T and Verizon in the two auctions. AT&T and Verizon were the two biggest winners in AWS-3. However, in the incentive auction, Verizon did not bid and AT&T was bidding to drop out of the auction completely after the first part of the auction. If AT&T and Verizon had behaved in a similar way in the AWS-3 auction, our estimate is that the AWS-3 auction would have generated net proceeds of less than $16.5 billion.

In the initial round of the auction, AT&T had the largest eligibility and largest commitment across all bidders. Its requested commitment peaked at $7,483,546,000 and its (processed) commitment peaked at $7,369,594,000, both after round 21 of stage 1. But then, between rounds 22-26 of stage 1, AT&T reduced its requested commitment to only $61,662,000. In the last round of stage 1 (round 27), AT&T was effectively bidding to drop out of the auction with bids only for Category 2 products and a requested commitment of $61,662,000. In round 2 of stage 1, AT&T further reduced its requested commitment to $10,821,000, again bidding only for Category 2 blocks. In all later rounds and until the conclusion of the clock phase, AT&T was consistently bidding to completing drop out of the auction, submitting zero quantity bids for all products.

A plausible explanation for the reversal in AT&T's bidding strategy is that the company's priorities changed at the time that the deal with Time Warner was made. Even though the deal was announced in October 2016, it has been reported that the first meeting between the CEOs of the two companies occurred on August 25,[5] that is, a few days before rounds 22-27 of stage 1 which occurred August 29-30. An alternative plausible explanation is that AT&T expected to win the FirstNet contract (which it was eventually awarded a few months later) to build a nationwide 700 MHz network for public safety, removing AT&T's need for additional low-band spectrum.

The auction rules were fairly robust to the reversal in AT&T's strategy, because a bid to reduce demand for a product could only be applied to the extent that it would not cause excess supply (no excess supply rule, see Section 4). Thus, AT&T ended stage 1 with a commitment of $4,187,955,820, even though its requested commitment at the end of stage 1 was only $61,662,000. Because the final stage rule was not met in stage 1, AT&T was able to reduce its commitment at the beginning of every new stage, as the supply of most products was reduced by one. Still, because every new stage was designed to be a continuation of the previous stage, AT&T's commitment could not be reduced to its requested commitment of $0. AT&T's commitment at the conclusion of the clock phase was $894,137,000.

Comparison with Auction 97 (AWS-3)

The AWS-3 Auction was a simultaneous multiple round auction which generated record net proceeds of $41.33 billion. In order to obtain an estimate of the AWS-3 net proceeds in a scenario where the behavior of AT&T and Verizon is along the lines of what it was in the incentive auction, we conducted a simulation.

[5] Code Names and Covert Meetings in AT&T's Courtship of Time Warner, The New York Times. Oct 24, 2016.

Using the publicly available bid information for AWS-3, we estimated a bidder's valuation for a license to be equal to maximum bid amount that the bidder submitted for that license. With the exception of Dish, we estimated a bidder's budget to be equal to its maximum commitment across all rounds. For Dish, we assumed that its net budget was equal to $10 billion. This is consistent with Dish's bidding, across its three affiliates, and the fact that Dish returned some of the licenses it won after the auction when the FCC determined that Dish should not have qualified for a bidding credit.[6]

We simulated the simultaneous multiple round auction assuming that in each round every bidder (except for Verizon and AT&T) bid on the most profitable set of licenses subject to its budget constraint, given the provisionally winning bids from the previous round. We assumed that Verizon did not participate in the auction. Finally, we assumed that AT&T stopped bidding in the auction during round 46 because this allowed for AT&T's final commitment to be very close to its final commitment in the incentive auction.

Table 1 shows, for each of the main bidders, the bidder's actual net payment in AWS-3 as well as its simulated net payment under the assumptions listed above. We find that if AT&T and Verizon had bid similarly to the incentive auction, the AWS-3 net payment would be reduced to $16.5 billion.

Table 1

Winning Bidder	Actual AWS-3 Net Payment	Simulated AWS-3 Net Payment
AT&T	$18,189,285,000	$ 896,501,000
Verizon	$10,430,017,000	$ 0
Dish / SNR Wireless / Northstar	$ 9,995,567,775	$ 9,989,998,000
U.S. Cellular / Advantage Spectrum	$ 338,304,000	$ 338,289,000
T-Mobile	$ 1,774,023,000	$ 3,499,989,000
2014 AWS Spectrum Bidco	$ 291,810,000	$ 817,271,000
Puerto Rico Telephone Company, Inc.	$ 170,901,300	$ 125,188,000
All other bidders	$ 139,765,250	$ 797,664,000
Total	**$41,329,673,325**	**$16,464,900,000**

[6] We note that this approach could be underestimating a bidder's valuation for a license in the event that the bidder won that license in AWS-3. An upper bound on revenues in the scenario where the level of participation for AT&T and Verizon is similar to the incentive auction can be calculated by summing the payments of AT&T and Verizon for the incentive auction with each other bidder's estimated budget for AWS-3. This approach yields an upper bound of $17.94 billion for the net proceeds that AWS-3 would have generated if the level of participation of AT&T and Verizon had been similar to the incentive auction. We note that $17.94 is still lower than the net revenue of the incentive auction.

3.4 Clearing Target Selection and Auction Duration

Bidding in the incentive auction commenced in March 2016 and concluded in March 2017. A total of 372 bidding rounds were conducted. Table 2 shows the number of rounds for each part of the auction.

Table 2

	Number of Reverse Auction Rounds	Number of Forward Auction Rounds
Initial Commitment	1	-
Stage 1	52	27
Stage 2	53	1
Stage 3	52	1
Stage 4	53	58
Assignment Phase	-	74
Total	211	161

One potential criticism of the incentive auction is that it lasted too long. If, after the initial commitment, the FCC had selected a clearing target of 84 MHz (instead of 126 MHz), it is very likely that the auction would have concluded after a single stage with significantly fewer rounds. As described below, an initial clearing target of 84 MHz would have had the additional advantage of reducing strategic bidding in the forward auction.

Once reverse auction bidders submitted their initial bids in the initial commitment period, the FCC used these bids as input to a clearing target optimization procedure which determined 126 MHz as the clearing target for stage 1. The clearing target was announced to bidders and the reverse auction for stage 1 was conducted. This determined that the cost for clearing 126 MHz was $84.6 billion. The clearing cost was made public before bidding for the forward auction of stage 1 commenced. This cost significantly exceeded all estimates on forward auction proceeds. As a result, forward auction bidders bid in stage 1 expecting that stage 1 would not be the final stage, which might have motivated forward auction bidders to use certain strategies to park eligibility during stage 1 (see Section 4.6).

The FCC had adopted procedures to conduct forward auction bidding in stage 1 regardless of the clearing cost in order to attempt to clear the maximum amount of spectrum. Similar considerations motivated the FCC's decision to select the next lowest clearing target in the 600 MHz band for any subsequent stage.

We now describe an alternative clearing target determination which would have likely resulted in selecting a lower clearing target for stage 1. The FCC would run the clearing target optimization procedure to determine a *provisional* initial clearing target. Reverse auction bidding for stage 1 would start with that clearing target (e.g. 126 MHz). However, if after any reverse auction round the cost for all provisionally winning reverse auction bidders exceeded some threshold, then the clearing target would be reduced and reverse auction bidding would continue bidding for the lower clearing target (without conducting a forward auction in between). The threshold would represent an upper bound on the expected forward auction revenues and could depend on the clearing target (e.g. the threshold could be $60 billion for 126 MHz and $50 for 114 MHz). This approach would have reduced the duration

of the auction but also possibly reduced the prevalence of parking in stage 1 of the forward auction (see Section 4.6 for details), because forward auction bidders would only be asked to bid for a clearing cost that could potentially be covered with forward auction proceeds.

An alternative approach for reducing the duration of the auction would be to only use the most spectrally efficient clearing targets. For instance, the FCC could have skipped the 108 MHz clearing target which has 8 forward auction licenses of 10 MHz each but requires clearing 108 MHz, an inefficiency that is a consequence of channel 37. This approach would have reduced the duration of the auction (by about 53 rounds) but would not have reduced the incentive for parking in the forward auction of stage 1.

4. The Clock Phase of the Forward Auction

In the clock phase of the forward auction, generic licenses were sold in 416 geographic areas (Partial Economic Areas, or PEAs). A product was the pairing of a PEA and a license category. Before the final stage rule was met, there could be at most two categories within a PEA: Category 1 (with impairment up to 15%) and Category 2 (with impairment between 15% and 50%). After the final stage rule was met, Category 1 could be split in at most two categories: reserved (C1-R) and unreserved (C1-U). Only a few PEAs had Category 2 blocks in the first three stages. In stage 4 (the stage that the final stage rule was met), there were no Category 2 blocks. As a result, in any given round there were at most two products for each PEA.

In the initial round of the auction (round 1 of stage 1), bidders indicated their demands for products at the opening prices. In every other round, there was a range of prices for each product. The highest price in the range is called the round's clock price for the product. A bidder could either maintain its demand at the round's clock price or request to change its demand at a price associated with the round. The clock prices of products with excess demand after a round were incremented.

A bid to change demand was considered a request and was applied during bid processing subject to the following constraints:

- No excess supply rule: If aggregate demand for the product was greater than or equal to supply, the product's aggregate demand after applying the bid could not drop below supply.
- Bidder's eligibility: The total number of bidding units associated with the bidder's demand after applying the bid could not exceed the bidder's eligibility in the round.

The no excess supply rule was applicable to bids for reducing demand. It guaranteed that (unless the product's supply increased in a later stage) no licenses of the product would remain unassigned if, after a round, the product's aggregate demand was greater than or equal to its supply. Furthermore, it guaranteed that proceeds increased after every round within a stage. Thus, if the final stage rule was met after a round, the stage could be declared to be the final stage.

The restriction on the bidder's eligibility was relevant for bids to increase demand if in the same round the bidder submitted one or more bids to reduce demand which could not be applied because of the no excess supply rule. The eligibility restriction limited the possibility of bid snipping.

A bid's price point is defined to be the percentage of the distance between the bid price and the lowest price in the round over the price range for the product.

Bids were processed in increasing order of price point, breaking ties with pseudorandom numbers. For instance, if a product had one unit of excess demand and two bidders submitted bids to reduce demand by one block each (and no other bidder bid to change its demand for that product), only the bid at the lowest price point would be processed.

4.1 Bidding Options and Bid Types

In the clock phase, forward auction bidders were not provided the full flexibility of package bidding as that would have been inconsistent with the no excess supply rule and the goal of meeting the final stage rule. Bidders could use four different types of bids to express how they wanted the bidding system to handle a bid to change demand if that bid could not be applied in full, either because of the no excess supply rule or because of the eligibility restriction.

The four bid types were: simple bids, switch bids, all-or-nothing bids (AON), and all-or-nothing bids with backstopping (AON+). We find that simple bids and switch bids were used extensively, the use all-or-nothing bids was more limited, and the use of all-or-nothing bids with backstopping was extremely limited.

In addition to the four bid types, bidders were also given the option of intra-round bidding which allowed them to submit bids at any price in the product's price range for the round. A bidder could either submit a simple bid to maintain its demand from the previous round at the current round's clock price or request to change its demand at a price associated with the round (using any type of bid). We find that most bids were either at the highest or at the lowest price of the round, and that most of the bids at intermediate prices were bids to increase demand.

A total of 138,691 bids were submitted in the clock phase of the forward auction.[7] In the initial round (round 1 of stage 1), bidders submitted 724 bids to indicate their demands at the opening prices. In later rounds, bidders submitted 120,402 bids for maintaining demand at the round's clock price and 17,565 bids for changing demand (i.e., either increasing or reducing demand).

For a meaningful comparison of the frequency of each bid type, Table 3 only considers bids to change demand after the initial round. The table also indicates whether the bid price was the lowest in the round's price range for the product (that is, at the 0% price point), an intermediate price (strictly between the lowest and the highest), or the highest price for the product in the round (the product's clock price for the round, or equivalently at the 100% price point).

Table 3

Bids to Change Demand After Initial Round	Price			All Prices
	Lowest	Intermediate	Highest	
Simple	5,915	2,295	2,295	10,505
Switch	4,282	785	1,872	6,939
AON	67	7	39	113
AON+	1	7	0	8

[7] The total number of bids (138,691) is smaller than the number of rows in the forward-bids file that is available on the PRS (145,638), because each switch bid and each all-or-nothing bid with backstopping is represented by two rows in that file. In the numbers provided in this document, each bid is counted once.

Total	10,265	3,094	4,206	17,565

Table 4 breaks down the bids to change demand into bids for increasing and reducing demand. Only simple and all-or-nothing bids are considered here, because by definition switch bids involve both a reduction and an increase and all-or-nothing bids with backstopping could only be placed for reductions in demand.

Table 4

	Bids to Reduce Demand	Bids to Increase Demand	Bids to Change Demand
Simple	6,631	3,874	10,505
AON	63	50	113

Simple Bids

A simple bid indicates a desired quantity for a product at a price. During bid processing, a simple bid was applied to the maximum extent possible subject to the no excess demand rule and the bidder's eligibility. For example, a simple bid to reduce a bidder's demand for a product from two blocks down to zero would be applied partially if there was only one unit of excess demand for the product.

Table 3 implies that 22% of all simple bids to change demand were at intermediate prices. However, the percentage of simple bids to reduce demand that were at intermediate prices is significantly lower. Table 5 provides a break-down between simple bids to reduce demand and simple bids to increase demand.

Table 5

Simple Bids[8] to:	Price			All Prices
	Lowest	Intermediate	Highest	
Reduce Demand	5,111	521	999	6,631
Increase Demand	804	1,774	1,296	3,874
Total	5,915	2,295	2,295	10,505

We observe that less than 8% of all simple bids to reduce demand were at intermediate prices (that is, with price points strictly between 0% and 100%). Furthermore, a lot of these bids were concentrated either around the lowest price or around the highest price for the round. In fact, only 112 (1.7%) of the simple bids to reduce demand had price points between 10% and 90%.

Intermediate prices were used more for bids to increase demand than for bids to reduce demand. The limited use of intra-round bidding for bids to reduce demand could be a consequence of the fact that the FCC used relatively low price increments to the set the clock prices of a new round.[9]

[8] After the initial round.

[9] The price increment was 5% in most rounds. The FCC increased the price increment to 10% in the latter parts of stage 1 and stage 4.

Bidders appear to be used intra-round bidding to prioritize bids to increase demand in the event that some of the bidder's reduction bids for the round could not be applied during bid processing. Since bids were processed in increasing order of price point, a bidder could prioritize one bid to increase demand over another bid to increase demand by associating a lower price point with the first bid.

A significant number of the bids to increase demand at intermediate prices were placed by Dish (1,204 out of 1,774). For example, in round 1 of stage 3 Dish bid to reduce its demand for 9 products and to increase its demand for 348 other products. Three of the reduction bids were not applied during bid processing because of the no excess supply rule. Thus, it was not possible to apply all of Dish's bids for increasing demand because that would have caused Dish to exceed its eligibility. The auction system attempted to apply bids with higher priority (according to Dish's bids) first. As a result, 241 of Dish's bids to increase to demand were applied and 107 were not applied.

Switch Bids

A switch bid is a request to move demand for up to n blocks in a given PEA from one category to another category. Thus, a switch bid could have been used in stages 1-3 to move demand from Category 1 to Category 2 (e.g., if the price for Category 1 was much higher than the price for Category 2) or from Category 2 to Category 1 (e.g., if the gap between the two prices was reduced in the previous round). Similarly, after the final stage was met and the reverse split was implemented, a switch bid could be used to request to move demand from the unreserved product to the reserved or vice versa.

A switch bid ensures that the bidder's total demand in the PEA remains the same. For example, if the bidder requested to switch one block from the unreserved to the reserved category but there was no excess demand for the unreserved product, the switch bid would not be applied and the bidder's demand would not change for either of these products.

If switch bids had not been available, in order to move its demand from an unreserved product to the reserved the bidder would submit two separate simple bids: one to reduce demand for the unreserved and one to increase demand for the reserved. Even if the reduction bid could not be applied because of the no excess supply rule, it could be plausible (depending on the bidder's eligibility and on what other bids it submitted in the round) that the bid to increase demand for the reserved product could be applied. In that case, the bidder's total demand in the PEA would increase. Switch bids eliminate this risk by linking the bid to increase demand for one category with the bid to reduce demand for the other category.

No switch bids were placed in stages 1-3. This was likely a consequence of the fact that the number of Category 2 products in stages 1-3 was very limited, and that bidders strategically used Category 2 products to park eligibility (see Section 4.6).

All the switch bids in the auction were made after the reserve split, that is, starting in round 3 of stage 4. Out of the 6,939 switch bids, 3,825 were requesting to move demand from the reserved to the unreserved product and 3,114 were requesting to move demand from the unreserved to the reserved product.

Some of the bidders expressed a clear preference for unreserved products, presumably because the reserved products were associated with certain restrictions. For instance, all of the switch bids submitted by Columbia Capital were to move demand from the reserved to the unreserved product.

However, most bidders submitted switch bids in both directions, depending on the relative prices of the reserved and the unreserved products in a PEA.

All-or-nothing Bids

All-or-nothing bids indicate a desired quantity of a product at a price, just like simple bids. However, all-or-nothing bids were either applied in full or not at all. These bids were available for a bidder that did not wish to place bids that could be applied partially. If an all-or-nothing bid to reduce demand could not be applied, the bidder would still hold its unreduced quantity potentially at a higher price.

As shown in Table 4, out of a total of 113 all-or-nothing bids 63 were for reducing demand and 50 were for increasing demand. The majority of all-or-nothing bids were submitted by AT&T and Spectrum Financial Partners. Only seven all-or-nothing bids were submitted by other bidders.

An intended use case for all-or-nothing bids was to allow a bidder to avoid holding one block when requesting to reduce its demand from two to zero blocks. There were 642 bids (all-or-nothing and simple) to reduce demand from two to zero blocks. 60 of these were all-or-nothing bids. 56 of these bids were by AT&T, three by Mark Twain Communications, and one by Financial Spectrum Partners. Thus, with the exception of AT&T, bidders hardly used all-or-nothing bids with the intention to avoid the possibility of holding one block after the round when submitting a bid to reduce demand from two to zero.

All-or-nothing bids were also used to manage eligibility. For example, in round 23 of stage 1, Chariton Valley Telephone Corporation submitted an all-or-nothing bid to reduce its demand in PEA 393 from five blocks to one block. This bid might appear counterintuitive at first, because one might expect that the bidder would try to avoid holding a quantity of one block. But in the same round, the bidder submitted a simple bid to increase PEA 180 by 1 unit. PEA 393 has 25 bidding units and PEA 180 has 100 bidding units. By submitting these bids the bidder ensured that a reduction in PEA 393 would be applied if and only if the reduction could be equal to 4 blocks in which case the increase in PEA 393 could be applied. If the bidder had submitted a simple bid for the reduction of 4 blocks, there would be a risk that the reduction bid would be applied only partially (not freeing up enough bidding units to apply the increase in PEA 180) in which case the bidder's eligibility would be reduced after the round.

All-or-nothing Bids with Backstopping

By submitting an all-or-nothing bid to reduce demand the bidder could ensure that it would not hold an undesirable quantity after the round, but the bidder could continue to hold the same (unreduced) demand at a higher price. An all-or-nothing bid with backstopping allowed the bidder to backstop the all-or-nothing bid at a higher price. Thus, if the price reached the specified backstop price, the bid would be applied in part (like a simple bid for the same quantity).

There were only 8 all-or-nothing bids with backstopping. 7 of these bids were submitted by the same bidder (Gold Spectrum, LLC) and were essentially the same bid repeated in four rounds. The backstop price of each of these bids was $1 higher than the all-or-nothing price. This implies that these all-or-nothing bids with backstopping had essentially the same effect as a simple bid at that price. We conclude that the use of all-or-nothing bids with backstopping was extremely limited.

The Absence of the Ability to Switch across Products (or PEAs)

In 2012, the FCC released the Incentive Auction Notice of Proposed Rulemaking, which included a high-level description of the proposed auction design in Appendix C (Milgrom, et al., 2012). That proposal had suggested that a bidder could specify a demand decrease for one product and a demand increase for another product (presumably in a different PEA) with the same bid, and that a bid could also involve multiple products or quantities.

The ability to submit switch bids involving products from different PEAs was not provided to bidders in the incentive auction, presumably on grounds of simplicity. That is, switch bids could only involve products in the same PEA. This had the potential of providing less flexibility to bidders. For instance, a bidder that wanted to reduce its demand by one block in each of two smaller PEAs in order to increase its demand by one block in a larger PEA would have to do so using only simple bids, and potentially lose some eligibility unintentionally if one of the two reduction bids could not be applied because of the no excess supply rule.

To find the extent to which bidders unintentionally lost eligibility (which they could have avoided if switch bids across products had been available), we look at all the times that a bidder's eligibility was reduced even though the bidder submitted activity that met or exceeded the required activity to maintain eligibility. There were only 18 such instances, and the reduction in the bidder's eligibility was less than 20% in each of these instances.[10] This suggests that the absence of the ability to submit switch bids across products was generally not harmful for bidders.

4.2 Reserved Spectrum

Concerns about market concentration have been large both in North America and the rest of the developed world. In 2015, at the time that the FCC adopted the reserved spectrum procedures for the incentive auction the market share of wireless subscriptions was 34% for AT&T, 33% for Verizon, 16% for Sprint and 16% for T-Mobile. Thus, AT&T and Verizon together held 67% of wireless subscriptions. Furthermore, as the FCC explained in the Incentive Auction Report and Order in 2014, two nationwide operators held 73% of sub-1-GHz spectrum, whereas the two smaller nationwide operators held 15% of the low-band spectrum.

The FCC thus adopted a mechanism to reserve spectrum for smaller bidders, but only after the final stage rule was met. The intention was to set aside spectrum for small bidders while ensuring that all bidders would pay their fair share for the final stage rule to be met. The reserved product was not identical to the unreserved product in a PEA, because the reserved product can only be sold to reserve-eligible bidders in the post-auction market. Thus, the unreserved price was expected to be greater than or equal to the reserved price in a PEA.

The only bidders that could not qualify to be reserve-eligible in all PEAs were AT&T and Verizon. However, even AT&T and Verizon were reserve-eligible in some PEAs. Furthermore, 69 PEAs had the property that all bidders could be reserve-eligible there. Because Verizon did not bid and AT&T lost interest in the auction during stage 1 (that is, before the final stage rule was met), essentially all of the bidders that were bidding after the spectrum reserve split were reserve-eligible.

[10] A bidder's eligibility was reduced by more than 10% in only 3 of these instances, and by more than 5% in only 6 of these instances.

That the spectrum reserve was nonbinding does not prove that it was unnecessary because bidders might have taken different actions if there were no reserved spectrum. It is equally likely that the policy was successful in making a foreclosure strategy impossible and in encouraging or emboldening participation by smaller bidders. We conclude that some form of spectrum cap and/or reservation continues to be very relevant for future auctions.

By design, the spectrum reserve was to be implemented during the bid processing of the round that the final stage rule was met. The fact that the spectrum reserve was implemented even though it was not binding appears to have had certain negative effects. First, bidding became more complex because the number of relevant products for each bidder was doubled, which potentially increased the likelihood of mistakes by bidders. Second, in some instances the reserved price was higher than the unreserved at the conclusion of the clock phase, even though the unreserved product is at least as valuable as the reserved.[11] Finally, implementation of the spectrum reserve may be increased the duration of the auction, because in some cases two rounds were required in order to increase the reserved and the unreserved price by one increment each even though only one round would have been required if the spectrum reserve had not been implemented and each PEA had a single product..

A possible improvement would have been to implement the spectrum reserve procedures only if there were at least two bidders bidding in the auction that could not qualify to be reserve-eligible in all PEAs. If this rule were in place, the spectrum reserve procedures would not have been implemented in the incentive auction because Verizon was not bidding. Such an approach would have the implication that bidders would after the final stage rule was met whether both AT&T and Verizon were bidding in the auction. This informational consequence would only occur towards the end of the auction and does not appear to have negative effects.

Another improvement to the reserve spectrum procedures would be to create a reserved product only if at least one of the bidders were not reserve-eligible in that PEA. Then, 69 of the PEAs would not have reserved products under any circumstances, because all bidders could be reserve-eligible there.

Finally, the auction design could have avoided the possibility of the reserved price being higher than the unreserved by adopting a different rule for incrementing prices. The price of a reserved (resp. unreserved) product in the auction was incremented if there was excess demand for that product, regardless of there was excess demand for the unreserved (resp. reserved) product in the that PEA. An alternative would be to increment prices in a way that guarantees that the price of the unreserved is always greater than or equal to the price of the reserved.

4.3 Final Stage Rule

The final stage rule consisted of two components. The first component required that forward proceeds from Category 1 products in the top-40 PEAs exceed a threshold in order to reflect competitive prices.

[11] In 140 PEAs, the unreserved price was larger than the reserved (in most cases 4-25% larger). In 195 PEAs, the unreserved price was equal to the reserved. In 81 PEAs, the unreserved price was smaller than the reserved; in most of those cases the reserved price was within one price increment of the unreserved, but there were a few cases where the two prices were two or more increments apart.

The second component required that forward auction proceeds net of bidding credits and impairment discounts be sufficient to cover the clearing cost for the stage, the estimated relocation costs for TV stations, and the cost of running the auction. Both components had to be satisfied in order for the final stage rule to be met. The final stage rule thus imposed two aggregate reserve price conditions: one based on what the FCC considered to be competitive prices, the other based on the clearing cost for the stage.

An aggregate reserve price condition requires bids across all bidders to be sufficiently high. Aggregate reserve prices can be problematic because of free riding. In particular, a bidder may refrain from increasing its bids because it expects or hopes that other bidders will bid sufficiently to meet the aggregate reserve price. This could reduce the chances that the aggregate reserve is met or equivalently increase the chances that auction fails.

An aggregate reserve price that attempts to ensure competitive prices can be problematic. A regulator can ensure a competitive price for each license by setting sufficiently high opening prices on a per-license basis. By setting an aggregate reserve price that is much higher than the sum of opening prices, the regulator faces the risk of setting the aggregate reserve too high. This risk is particularly relevant if competitive prices are defined based on previous auctions, since market conditions might have changed.

It appears that the threshold for the first component of the final stage rule had been set too high. For stage 4, the first component required that the average price per MHz-pop for Category 1 products in the top-40 PEAs be greater than or equal to the target of $1.25. The average price per MHz-pop for Category 1 products in the top-40 PEAs at the end of the clock phase was $1.27, just two cents above the $1.25 threshold.

The auction experienced a near miss in stage 4 because of the $1.25 threshold for the first component. In stage 4, the cost of the reverse auction dropped significantly to $10.05 billion (from $40.3 billion in stage 3) and, as a result, the second component of the final stage rule was met. The first component of the final stage rule was not met after round 1 of stage 4. In round 2 of stage 4, T-Mobile increased its demand in Los Angeles from 2 to 3 blocks causing the final stage rule to be met. If T-Mobile had not increased its demand in Los Angeles, it is plausible that the final stage rule would not have been met in stage 4 (see also discussion in Section 4.4) which would have caused the auction to proceed to stage 5 with a lower clearing target. In this hypothetical scenario, the final stage rule would have been met in stage 5 resulting in inefficiently little clearing.

The second component of the final stage rule (the cost component) was mandated by the Spectrum Act and represented a reasonable aggregate reserve price as it reflected the value of the spectrum to TV stations. In the first three stages, the forward auction proceeds were significantly lower than the cost of the reverse auction at that clearing target. Thus, the second component of the final stage rule (the cost component) was not met. The second component was met in stage 4.

The FCC has imposed aggregate spectrum reserve prices in previous auctions. In some cases, the aggregate reserve price was along the lines of the second component of the final stage rule in the sense that it was required to cover relocation costs. In Auction 66 (AWS-1) in 2006, the FCC established an aggregate reserve price for all AWS-1 licenses in order to implement a Congressional mandate to recover estimated relocation costs of government incumbent operators in the band. In Auction 97

(AWS-3), the FCC used an aggregate reserve price for each sub-band based on the total estimated relocation or sharing costs contained in the Federal agencies' transition plans.

The first component of the final stage rule is along the lines of the aggregate reserve in Auction 73 in 2008, where the FCC used a separate aggregate reserve for each block to promote the recovery of a portion of the value of the public spectrum resource, an approach that has been criticized (Cramton, et al., 2007). The aggregate reserve price for each block was set based on a variety of factors, such as the characteristics of the band and auction prices of other recently auctioned licenses.

In the context of a combinatorial clock auction it has been shown that bundle reserve prices, where the payment for a bundle of goods is required to be at least the sum of the reserve prices for the items in the bundle, can be problematic (Ausubel & Baranov, 2016). Aggregating the reserve price across bidders exacerbates the problem because of free riding.

4.4 Extended Round

By design, an extended round would be triggered if, after the bid processing of a round, (i) the final stage rule was not met, (ii) there was no excess demand in the top-40 PEAs, and (iii) the final stage rule would be met under a scenario where the prices of the top-40 PEAs increased by 20% and all bidders maintained their demand. The extended round increment would then be set at 133% of the increment required to meet the final stage rule, which would make it possible for the final stage rule to be met even if some bidders did not maintain their demand at the extended round clock prices.

The extended round was designed to give bidders a last chance to meet the final stage rule in the current stage and thus avoid having the clearing target reduced.

No extended round was triggered in the auction. In the first three stages, the forward auction proceeds at the time that there was no excess demand in the top-40 PEAs were significantly lower than the cost of the reverse auction at that clearing target. Thus, the auction proceeded to the following stage, with a lower clearing target.

In stage 4, the cost of the reverse auction dropped significantly to $10.05 billion (from $40.3 billion in stage 3) and, as a result, the cost component of the final stage rule was met. The first component of the final stage rule was not met after round 1 of stage 4. At that time, there was excess demand for some of the top-40 PEAs, so the auction proceeded to round 2 of stage 4 (without triggering an extended round). In round 2 of stage 4, the final stage rule was met, eliminating the need for an extended round.

In the hypothetical scenario that T-Mobile would not increase its demand in Los Angeles in round 2 of stage 4, it is plausible that an extended round would have been triggered in stage 4. Whether the final stage rule would have been met in such an extended round depends on whether the bidders, collectively, would have wanted the extended round to succeed. Because AT&T was trying to get out of a number of top-40 PEAs in stage 4, it is plausible that the extended round would have failed in which case the auction would have proceeded to stage 5.

When the extended round was designed it had not been anticipated that a bidder holding demand in multiple top-40 PEAs would be trying to reduce its demand at a price that it previously accepted. It had been assumed that bidders would only consider reducing their demand at higher prices and thus it appeared reasonable to overshoot the extended round clock prices by 33%. However, that would not

be enough for bidders to meet the final stage rule in a scenario where a large bidder is attempting to drop out of the auction. In such a scenario, it would be more appropriate to overshoot by more than 33%.

4.5 Unassigned Licenses

In the final stage of the auction, there was a supply of 7 blocks in each of the 416 PEAs, corresponding to a total of 2,912 licenses. 2,776 of these licenses were sold in the auction and 136 remained unassigned.

In the top 40 PEAs, there was one unassigned license in Los Angeles (PEA002) and two in San Diego (PEA018). That is, only 6 out of 7 blocks sold in Los Angeles and only 5 out of 7 in San Diego. This occurred despite the fact that, for each of these PEAs, during stage 1 aggregate demand at the price that ended up being the final clock price exceeded 7 blocks.

The unassigned licenses in Los Angeles and San Diego were caused both by insufficient demand for those PEAs in stage 4 and by the fact that the supply in each of those PEAs was larger in stage 4 than it was in stage 1. The latter was a consequence of an agreement between the FCC and the Mexican communications regulator, the *Instituto Federal de Telecomunicaciones* (IFT), whereby the two agencies agreed to a channel plan which identified future channels for operating Mexican stations along the border with the United States that essentially moved all Mexican channels below channel 37 (with a few exceptions that are listed in the coordination documents). However, the auction design was done before the agreement between the FCC and IFT and did not contemplate the possibility that supply could increase between stages.

4.6 Strategic Bidding

The main attempts for bidding manipulation that we observed in the clock phase of the forward auction had to do with parking eligibility with the goal of delaying bids that reflected bidders' true preferences. Bidders tended to park in the top PEAs (that is, the PEAs with the most bidding units) and in Category 2 products. Bidders were able to park more eligibility in the top PEAs than in Category 2, but, as explained below, that strategy was also more risky.

The prevalence of parking in the first stage of the auction might have been due to the fact that bidders did not expect stage 1 to be the final stage. At the time that bidding in the forward auction started, bidders knew that the clearing cost was $86.4 billion and thus the forward auction net proceeds would need to be at least $88.6 billion for the final stage rule to be met. This number significantly exceeded expectations on net forward auction proceeds, thus bidders were pretty confident that the final stage rule would not be met in stage 1. Nevertheless, because a new stage was designed to be a continuation of the previous stage, bids made in stage 1 for Category 1 products could not necessarily be reversed in later stages.

Parking Eligibility in Top PEAs

In the first rounds of stage 1, the aggregate demand for some of the top PEAs, such as New York (PEA001) and Los Angeles (PEA002), was extremely high, indicating that bidders were parking their eligibility. For example, in the initial round, two of the bidders bid for 10 blocks in New York (PEA001) each, a significantly larger quantity than any bidder was expected to be interested in.

We next provide some evidence that bidders were not expressing their true demands in stage 1 by considering the bidding behavior in Los Angeles (PEA002). In stage 1, the supply for Los Angeles was equal to 5 blocks. In round 25 of stage 1, bidders indicated aggregate demand for 8 blocks at the clock price of $388,489,000. However, in the following rounds of stage 1, bidders requested only 4 blocks in aggregate at the same price ($388,489,000). Because of the no excess supply rule, the aggregate processed demand remained at 5 blocks (that is, equal to the supply for stage 1). In later stages, the supply of blocks in Los Angeles increased due to an international agreement with Mexico. During stage 4, the aggregate demand for Los Angeles increased to 6 blocks; however, the supply was equal to 7 and, as a result, there was one unsold block in Los Angeles even though the final clock price in Los Angeles was $388,489,000 (that is, the price did not increase since round 25 of stage 1).

The information provided in Table 6 also suggests that bidders were parking eligibility during stage 1. The table shows, for each of the six largest winners in the auction, the number of PEAs bid in the initial round, the number of PEAs bid in round 20 of stage 1 (i.e., just before AT&T reversed its strategy), and the number of PEAs where the bidder had winnings.

Table 6

Bidder	Number of PEAs Bid in Initial Round	Number of PEAs Bid in Round 20 of Stage 1	Number of PEAs with Winnings
T-Mobile	409	409	414
Dish	54	59	416
Comcast	6	38	72
Columbia Capital	3	3	6
AT&T	13	67	18
Bluewater	3	66	64

With the exception of T-Mobile, the ratio of the number of PEAs bid by the bidder in the initial round over the number of PEAs where the bidder won blocks is generally small for these bidders. This suggests that most of these bidders were attempting to park eligibility early in the auction without revealing their true preferences.

However, this strategy was risky. Even though the final stage rule was not met in stage 1 (as expected), the aggregate demand in New York could only be reduced by one block at the beginning of a new stage, because the supply of blocks in New York was reduced by one block and the no excess supply rule implied that aggregate demand could not be reduced below supply. When stage 1 ended, multiple bidders were holding more blocks in New York than they were interested in and tried to reduce demand at the beginning of every new stage. Some of the bids could not be applied[12] and, as a result, one of the bidders appears to have won more blocks in New York than it wanted.

[12] Because these reduction bids were at the 0% price point, the auction system used pseudorandom numbers to select which bid to apply.

Parking Eligibility in Category 2 Products

Category 2 blocks were blocks with an impairment level between 15% and 50%. Because the impairment level of the Category 2 blocks offered in a given PEA could vary significantly from one stage to the next (e.g. it could be 16% in stage 1 and 45% in stage 2), the FCC decided to not apply the no excess supply rule for Category 2 products at the beginning of a new stage. That is, a bid to reduce demand for a Category 2 product would be applied in full during the bid processing of the first round of a new stage, even if there were no excess demand for that product.

However, the reverse auction cost in stage 1 was so high that bidders expected that stage 1 would not be the final stage. Some of the bidders then realized that they could park some eligibility in Category 2 products with the intention of using that eligibility elsewhere in a later stage. The extent of parking in Category 2 blocks was limited though, because only a small number of Category 2 blocks was available and the number of bidding units associated with those products was not significant.

Bidders also parked some eligibility in Category 2 blocks during stage 2 and stage 3. However, this was even more limited than in stage 1, because the number of Category 2 blocks available was even smaller.

5. The Assignment Phase

The clock phase of the forward auction determined the number of generic blocks for each product awarded to each bidder. The purpose of the assignment phase was to map generic winnings into actual physical frequencies. The assignment phase was organized as a sequential sealed-bid auction where bidders bid for specific frequency assignments in each region (PEA) independently, one after another. The regions were sorted in a descending order of population. Participating bidders were informed about their final assignments in each region before bidding for the next region.

Bidding in the assignment phase was optional since winners of the generic blocks were guaranteed some assignment. For example, a bidder who was indifferent between all of its possible assignment options did not need to bid at all. In fact, 18 out of 50 clock phase winners chose not to bid in the assignment phase.

Given the total of 416 PEAs, processing regions one by one would have taken a long time. To facilitate a faster process and to produce a less fragmented assignment, FCC introduced two features: grouping and parallel bidding.

5.1 Grouping of PEAs and Parallel Bidding

Grouping feature allowed FCC to merge several regions into a larger aggregation for a simultaneous assignment. In particular, regions were grouped when the same assignment could be made in all of the regions in the group for each clock phase winner.[13]

The rationale behind grouping regions is as follows. When designing a spectrum auction, the regulator faces an important “product design” decision about how finely to partition the country into licenses. If

[13] A necessary condition for grouping was for each region in a group to have the same clock phase winners and winnings.

the partition is too coarse, small bidders could be seriously harmed because they are forced to bid on large regions that do not match their footprint. In contrast, if the partition is too fine, the product space, complexity and duration of the auction are unnecessarily increased. In general, there is a preference of erring on the finer side, because a finer partition, albeit the auction complexity, does not preclude bidding and lets bidders to assemble larger aggregations. In this context, grouping of regions in the assignment stage can be viewed as a systematic way of undoing any "too fine" partition errors by ensuring that each bidder is assigned the same frequencies in each of the regions in a group, increasing the efficiency of the aggregate assignment.

The parallel bidding feature was introduced to further shorten the duration of the assignment phase. After all major markets (top-40 PEAs) have been assigned, the FCC started to run parallel independent assignment auctions in regions that are far away from each other geographically (regions from different REAGs[14], one region from each predefined REAG). The parallel bidding reduced the amount of feedback that bidders would have received in a fully sequential process, but, at the same time, the value of such feedback (specific frequencies a bidder won in a region far away) appear to be unimportant.

The grouping of PEAs and the parallel bidding for REAGs seemed to work well. Many the groupings consisted of PEAs that were close geographically and thus facilitated the goal of assigning a bidder the same set of licenses in a region. Some of the groupings consisted of a large number of PEAs, especially during parallel bidding. For instance, in round 27, in REAG 3, the grouping consisted of 22 PEAs out of which 21 were in Wisconsin and Iowa. As another example, in round 27, in REAG 5, the grouping consisted of 16 PEAs out of which 14 were in Texas. There were, however, some exceptions, particularly in the top-40 PEAs. The most notable exception is probably the grouping of assignment round 5 which consisted of Orlando (PEA013), Portland (PEA019), Pittsburgh (PEA023), and Fresno (PEA034).

If PEAs had not been grouped and no parallel bidding had been conducted, an assignment round would have been required for each PEA and thus the number of required assignment rounds would have been 416. In comparison, the actual assignment phase consisted of only 74 assignment rounds.

5.2 Vickrey Payment Rule

For the assignment phase, the FCC adopted a VCG payment rule in order to encourage bidders to bid according to their true values. While the VCG rule has good incentive properties, it also known to have some undesirable features (see Ausubel and Milgrom (2006)). The culprit of the problems traditionally associated with the VCG rule is that the payments are not guaranteed to be in the core (i.e., bidders may pay less than the amount offered by their competitors). This shortcoming has motivated the use of core-selecting payment rules[15] in recent spectrum auctions in the UK, Canada and Australia.

The bidding data from the assignment stage of the Incentive auction presents us with a unique opportunity to evaluate the extent of the core problem of the VCG. Treating bids as true values, we find that 38 instances (out of 228 total assignment auctions held by FCC) where VCG payments are too low to be in the core. The cumulative revenue shortage across all assignment auctions is $4,411,699. That is, assuming the same set of bids, if the FCC had used a core-selecting payment rule, the assignment phase gross revenues would have been $140,342,331 instead of $135,930,632, that is, about 3.25% higher.

[14] REAG stands for Regional Economic Area Grouping. FCC predefined 6 REAGs.
[15] See Day and Raghavan (2007), Day and Milgrom (2008), and Day and Cramton (2012).

However, given that core-selecting auctions are not incentive-compatible in general (i.e., bidders can gain by bidding in a non-truthful manner), any implications of this comparison are quite limited.

Ausubel and Milgrom (2006) provided a well-known example of the VCG rule where winners pay zero even though losers bid competitively. Their theoretical example was driven by the same core problem of the VCG rule described above. Using the bidding data from the assignment phase, we find 3 instances (out of 38 instances with core violations) where bidders literally pay zero. To the best of our knowledge, it is the first time that the "zero revenue" irregularity of the VCG rule has been observed in practice. We provide details for one of these instances below.

Example of a "Zero Revenue" outcome – PEA 224 (Assignment Round 37 in REAG 3)

The assignment for seven blocks ABCDEFG in PEA224 (De Kalb, IL and Kenosha, WI) has been made in round 37 (one of the rounds with parallel bidding). In this region, Dish Networks won one generic block, while T-Mobile and U.S. Cellular won three generic blocks each. Assignment stage bids submitted by bidders for this region are listed in Table 7.[16]

Table 7: Assignment Bids for PEA224

Bidders	*Dish Network*	*T-Mobile*	*U.S. Cellular*
Bids:	**Bid $0 m for A**	Bid $0.4 m for ABC	Bid $0 m for ABC
	Bid $0 m for D	**Bid $3.5 m for BCD**	Bid $0 m for BCD
	Bid $0.237 m for G	Bid $0.01 m for DEF	Bid $4.0 m for DEF
		Bid $0 m for EFG	**Bid $8.7 m for EFG**
VCG payment:	0	0	0

The mechanics of the VCG payment calculations for this example is as follows. First, observe that U.S. Cellular's bid of $8.7 m for blocks EFG prevented Dish Network from winning block G. Second, T-Mobile's bid of $3.5 m for blocks BCD prevented Dish Network from winning block G because winning it would have forced U.S. Cellular to a noncontiguous assignment. Therefore, in the absence of either T-Mobile or U.S. Cellular, Dish Network would still have won block A. As a result, their VCG payments equaled zero! Clearly, T-Mobile and U.S. Cellular paid to the FCC less than the amount offered by Dish Network for block G. In contrast, a core-selecting rule that was used in other spectrum auctions would have increased the joint payment of T-Mobile and U.S. Cellular to $0.237 m and avoided the "zero revenue" outcome.

"Zero revenue" outcomes and the high incidence of outcomes that lie outside of the core raises an important question about the FCC's motivation for adopting the VCG rule instead of a core-selecting rule. One potential explanation is as follows. The VCG rule sometimes allows bidders to acquire desired frequencies cheaply. Expecting a relatively cheap assignment round, bidders can afford to bid higher for the generic blocks in the allocation phase. For the incentive auction, only revenues from the allocation phase counted towards meeting the final stage rule. So there was a desire to pack as much revenues into the allocation phase as possible, resulting in a decision to make the prices in the assignment phase as low as possible.

[16] Infeasible bids (ones that could not win) are omitted here.

5.3 Hiding Feasible Assignment Options

With adoption of the VCG rule, there was a need for mitigating measures to make it harder for bidders to manipulate the assignment phase and damage their opponents. Another innovation made by the FCC for the assignment round was intended to accomplish these objectives by hiding the information about the winnings of the other bidders.

In a typical assignment round,[17] the auctioneer accepts bids on assignment options that are maximally consistent with assigning contiguous spectrum to all winners. For example, suppose that there are two winners in a given region who won 4 and 3 generic blocks respectively. Then the winner of 4 blocks would be presented with two assignment options: ABCD and DEFG. Note that this bidder would be able to infer that there is only one other winner in this region who won 3 generic blocks. To hide this information, the auctioneer can present bidders with all options that are consistent with contiguous assignments only for themselves. In our example, the winner of 4 blocks would be presented with four assignment options: ABCD, BCDE, CDEF and DEFG. Note that this policy still allows bidders to express all their tradeoffs among assignments for their own use. The extra assignment options does not affect the outcome of the assignment round since they are infeasible in reality, but they do allow to obscure potentially sensitive information from bidders. For example, the winner of 4 blocks, when presented with 4 assignment options, cannot infer how many other winners are there (can be 1, 2, or even 3) or how many generic blocks each opponent won. The one small drawback of this policy is a waste of resources borne by bidders to evaluate such infeasible options.

This innovation is probably not as critical in an auction with a core-selecting rule. However, in an auction with the VCG rule, where bidders can target their vulnerable opponents cheaply or even for free, this policy appears to be easily justified.

Once a decision was made for there to be no core adjustment, one tries to think of mitigating measures to make it harder for bidders to manipulate the assignment phase. It was in this context that we decided to output all (feasible and infeasible) options to bidders. This innovation is probably not as critical in an auction where a core adjustment is being used, but it still seems like a useful innovation. The key justification may be that we want to make it harder for firms to do damage to their opponents. Outputting all options does make it harder to target vulnerable opponents, while still allowing the bidder to express its tradeoffs among assignments for its own use.

6. Conclusion

[To be added]

[17] For example, see the rules for the Canadian 700 MHz and 2500 MHz auctions.

References

Ausubel, L. M. & Baranov, O., 2017. A practical guide to the combinatorial clock auction. *The Economic Journal,* forthcoming.

Ausubel, L. M. et al., 2014. Demand reduction and inefficiency in multi-unit auctions. *Review of Economic Studies,* Volume 81, pp. 1366-1400.

Ausubel, L. M. & Milgrom, P., 2006. The lovely but lonely Vickrey auction. In: *Combinatorial auctions.* Cambridge, MA: MIT Press.

Binmore, K. and P. Klemperer (2002), "The Biggest Auction Ever: The Sale of the British 3G Telecom Licences," *The Economic Journal*, 112, C74-C96.

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

Cramton, P., Rosston, G., Skrzypacz, A. & Wilson, R., 2007. *Comments on the FCC Wireless Bureau's Proposed Competitive Bidding Procedures for Auction 73,* s.l.: s.n.

Day, R. & Milgrom, P., 2008. Core-selecting package auctions 36, no. 3. *international Journal of game Theory,* 36(3), pp. 393-407.

Day, R. W. & Cramton, P., 2012. Quadratic core-selecting payment rules for combinatorial auctions. *Operations Research,* pp. 588-603.

Day, R. W. & Raghavan, S., 2007. Fair payments for efficient allocations in public sector combinatorial auctions. *Management Science,* 53(9), pp. 1389-1406.

Doraszelski, U., Seim, K., Sinkinson, M. & Wang, P., 2016. *Ownership concentration and strategic supply reduction,* s.l.: s.n.

Jehiel, P. and B. Moldovanu (2001), "The European UMTS/IMT-2000 License Auctions," Working Paper 1–20, University of Mannheim.

Loertscher, S., Marx, L. M. & Wilkening, T., 2015. A Long Way Coming: Designing Centralized Markets with Privately Informed Buyers and Sellers. *Journal of Economic Literature,* 53(4), pp. 857-97.

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

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

Milgrom, P., Ausubel, L., Levin, J. & Segal, I. (2012), "Incentive Auction Rules Option and Discussion," Appendix C to *Expanding the Economic and Innovation Opportunities of Spectrum Through Incentive Auctions*, Notice of Proposed Rulemaking, Docket No. 12-268, 27 FCC Rcd 12357 (2012), https://apps.fcc.gov/edocs_public/attachmatch/FCC-12-118A2.pdf.

Milgrom, P., L. Ausubel, J. Levin and I. Segal (2013), "Auctionomics/Power Auctions Option for Forward Auction," Supplement to Appendix C of the Notice of Proposed Rulemaking, https://ecfsapi.fcc.gov/file/7022116356.pdf.

Milgrom, P. & Segal, I., 2015. Deferred-Acceptance Auctions and Radio Spectrum Reallocation.

Appendix: Summary of Results

The auction concluded after four stages. Table 8 shows the clearing cost and the forward auction proceeds for each stage.

Table 8

Stage	Clearing Target	Licensed Spectrum	Clearing Cost	Gross Proceeds
1	126 MHz	100 MHz (10 blocks)	$86,422,558,704	$23,108,037,900
2	114 MHz	90 MHz (9 blocks)	$54,586,032,836	$21,519,907,210
3	108 MHz	80 MHz (8 blocks)	$40,313,164,425	$19,676,240,520
4	84 MHz	70 MHz (7 blocks)	$10,054,676,822	$19,632,506,746

After the assignment phase, the gross proceeds increased to $19,768,437,378. Only about 0.7% of the proceeds came from the assignment phase, suggestive that all seven blocks were viewed by bidders as nearly equivalent and that it was appropriate to treat all blocks as generic spectrum in the clock phase.

Out of 62 bidders that were qualified by the FCC to bid in the auction, there were 50 winning bidders. Table 9 provides the winners with the largest commitments at the conclusion of the auction.

Table 9

Winner	Gross Commitment	Net Commitment (if different than gross)
T-Mobile	7,993,361,993	
Dish Network	6,211,154,496	
Comcast	1,724,877,376	
Columbia Capital	1,008,704,549	858,704,549
AT&T	910,202,302	
Bluewater	718,323,225	568,323,225
New Level	348,878,183	296,546,456
U.S. Cellular	328,661,977	