

wireless bidder; station N will be paid $v_W - v_S$, station S will be paid $v_W - v_N$, and the winning buyer will pay v_W . Note that, when a channel is cleared, the payments deficit is $v_W - v_N - v_S > 0$, which is equal to the gain from the exchange.

By the celebrated Revenue Equivalence Theorem (8), this large deficit is a feature of not just the Vickrey auction but any fully efficient economic mechanism when agents are privately informed about their values, have property rights as described above, and behave to advance their self-interest. Conversely, according to the Myerson–Satterthwaite theorem (9), the inefficiency obtained in a deficit-free mechanism under those conditions has to be quite large. The situation in the real incentive auction is in one respect much worse than that in our simple example: to clear even a minimal amount of usable nationwide spectrum, hundreds of stations would need to be acquired. In that situation, if the auction mechanism is never allowed to incur any deficit, then it must have a vanishingly small probability of clearing the channel (10).

Fortunately, the magnitude of the deficit problem depends on the participants' property rights, which Congress was able to specify. The key to the analysis is to focus on how the price p_j for any bidder j is computed from the bids by other participants. If participant j 's VCG price p_j is a nondecreasing function of the bid \hat{v}_j of some other participant j' , then we say that j' is a "substitute" for j , whereas if p_j is a nonincreasing function of \hat{v}_j , then we say that j' is a "complement" for j . More aggressive bidding by bidders who are substitutes for bidder j will reduce bidder j 's payoff in the VCG auction: it will reduce p_j if j is a seller or raise p_j if j is a buyer. Thus, specifying property rights to create more substitutes for j reduces the VCG deficit. Specifying rights so that there are fewer complements has a similar deficit-reducing effect.

How can changing property rights make television stations into substitutes for each other? In our example above, each station was assumed to have the right to broadcast without interference on its current channel. No other station could substitute for it to enable clearing of the channel. The Middle Class Tax Relief and Job Creation Act of 2012, however, specified property rights differently. Each station that qualifies for protection has only the right to interference-free coverage in its service area on some channel. One consequence of this is that the price that a station can demand is reduced. A station that demands too high a price does not need to be bought: it can be assigned to continue broadcasting using a different channel. This "retuning" possibility promotes competition among stations on different channels and makes them substitutes.

In our example, there is no currently unused channel onto which stations N and S can be retuned: channel 3 is already occupied by station C . However, with the rights as defined by the 2012 act, if stations N and S demand prices that are too high, the auctioneer could instead purchase rights from station C . It could then retune N and S to broadcast to channel 3. In a VCG auction with these rights, the Vickrey price for station N is $p_N = \min(\hat{v}_C, \hat{v}_W) - \hat{v}_S$. With these property rights, C is a substitute for N : lower bids by C reduce N 's VCG price.

Determining How Much Spectrum to Clear. Given the large costs of television to wireless interference and the benefits of handset standardization, the FCC decided that (with limited exceptions) the same channels should be reassigned to wireless use in every part of the United States, but how many channels should that be? In economics textbooks, which treat markets for a homogeneous good, the efficient quantity of the good is one for which there is a market-clearing price at which quantities supplied and demanded are equal. For the incentive auction, however, there

is no homogeneous good and no single price. Every station's broadcast license covers a different population, every wireless license is distinct, and broadcast and wireless licenses are different from one another.

Ensuring That the Auction Has No Deficit. As we have seen, the redefinition of property rights reduces the deficit of the VCG mechanism, but it may not eliminate that deficit completely. In our simple example, if $\min(v_A + v_B, v_C) < v_W < \max\{v_A + v_B, v_C\}$, then the mechanism will implement the efficient outcome of clearing one channel, but total payments to broadcasters will still exceed revenues from the sale of wireless licenses.

In fact, there is necessarily a deficit even in the simple case of a homogeneous-good market, in which each seller has one unit of the good for sale and each buyer seeks to buy one unit. In that case, the set F of feasible allocations is made up of the zero–one vectors satisfying the constraint $\sum_j x_j = K$, where K is the number of sellers. In this example, each buyer's VCG price is equal to the highest value of an agent who leaves the market without a good, and each seller's VCG price is equal to the lowest value of an agent who leaves the market with a good; therefore, the sellers always pay (weakly) less than the buyers receive. However, the gap between the buyers' and sellers' VCG prices is likely to be quite small when there are many buyers and sellers. Furthermore, McAfee (11) suggested an alternative mechanism to VCG for large homogeneous-good markets that is deficit-free, strategy-proof, and nearly efficient: it sacrifices at most one valuable trade (the least valuable one) to ensure that the buyers' price exceeds the sellers' price.

The FCC's problem is more complicated than the homogeneous-good market described above. Stations in different areas are sometimes substitutes and sometimes complements, buyers with different interests are sometimes substitutes and sometimes complements, and buyers and sellers are most often (but not always) complements. In such complicated settings, the VCG auction typically suffers a deficit (12). Nevertheless, the idea by McAfee (11) can be extended to design a deficit-free incentive auction that determines the number of channels that the auction seeks to clear. Eliminating the deficit is accomplished by reducing the number of channels to be cleared below the socially efficient number. The resulting inefficiency need not be large, however, provided that there is sufficient substitutability among the spectrum sellers and among the spectrum buyers.

Broadcasters' Property Rights and Computational Challenges.

According to economic theory, efficient trading is promoted when property rights are defined so that nobody except the owner cares about the owner's identity. Such a definition ensures that transfers of property do not impose adverse externalities on third parties, and therefore, any mutually beneficial exchange must increase efficiency. For the incentive auction, the corresponding idea is that retuning any station to a different channel should not affect other channels (that is, it should impose no more than minimal interference for other stations).

The way that stations' noninterference rights were specified, however, affects the computationally feasibility of the auction. It takes days of computer time to quantify the number of customers affected by interference under a single assignment of channels to stations. With most of the 2,993 US and Canadian stations needed to be assigned postauction into 29 or more remaining television channels, the number of possible assignments could be as high as 29^{2993} , which is, roughly, a 4,300-digit number (to put this in perspective, the number of atoms in the universe has roughly

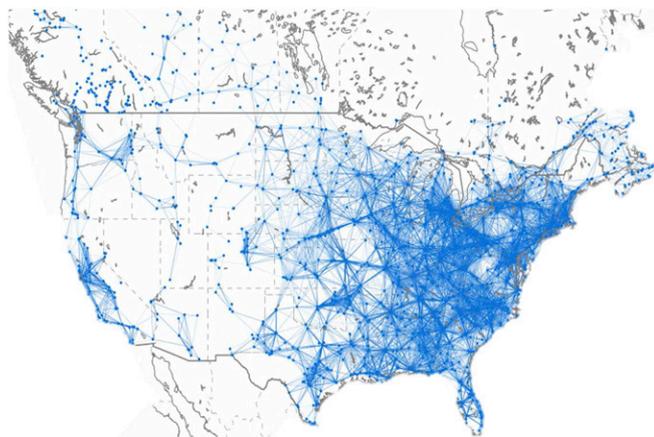


Fig. 1. Visualizing the incentive auction's 2.7 million pairwise interference constraints.

80 digits). Thus, it seems to be extremely challenging to indicate whether any particular set of stations could be feasibly assigned to the remaining television spectrum while exactly preserving each station's preauction coverage.

To keep the auction computations tractable, the FCC decided to treat interference differently. A station j would be deemed to suffer only minimal interference if no other single station interferes with more than 0.5% of station j 's preauction audience. In principle, this definition allows that the aggregate interference from neighboring stations could be substantially larger than 0.5%, but simulations found that it was never too much larger. This specification of acceptable interference simplified the auction by making it possible to verify the feasibility of an assignment by checking a few million pairwise channel-specific interference constraints of the form "station A cannot be assigned to channel X if station B is assigned to channel Y" (where X and Y could be the same or adjacent channels) as illustrated in Fig. 1. These pairwise constraints were computed in advance of the auction.

Even with this simplification, however, the problem of determining whether a given set of stations can be feasibly assigned into a given set of channels (henceforth, the "feasibility checking problem") is NP-complete (it generalizes the well-known "graph coloring problem"). Applying the widely believed hypothesis from complexity theory that $P \neq NP$,* any feasibility checking algorithm must have worst case running time that grows exponentially in the size of its input, meaning that it will fail to terminate within a reasonable amount of time for at least some inputs.

To summarize, among the relevant economic analyses are the Coasian analysis, which emphasizes that groups have incentives to find and coordinate on a good plan, and the Hayekian analysis, which emphasizes that the easiest way to find a good plan is sometimes to find market-clearing prices by a decentralized procedure. For the incentive auction, however, finding a good plan (including how many channels to clear, which stations to retune, and how to retune them in a way that satisfies millions of interference constraints) and the clearing prices while still enforcing all of the interference constraints is a computationally hard problem, which requires a coordinated solution involving very many

*One implication of denying the hypothesis that $P \neq NP$ is that there must exist a single general algorithm capable of defeating virtually all of the encryption systems in common use.

parties. Thus, a central authority like the FCC has an indispensable coordinating role. It must ensure that broadcasters' individual plans are mutually compatible and that, in combination, they are feasible and promote an approximately efficient allocation.

Market Design

A Last Look at VCG. As described above, the VCG auction design is the unique sealed bid auction that provides incentives for truthful bidding, implements efficient allocations, and involves no payments to or from losing bidders. However, there are several reasons why the unmodified VCG auction would be unsuitable for the incentive auction. First, as noted above, it generally runs a budget deficit. Second and perhaps more importantly, it would require solving the computationally challenging problem of finding a value-maximizing allocation subject to the millions of interference constraints as well as a separate instance of this problem for each winning bidder, assuming this bidder's nonparticipation, to calculate the bidder's VCG price. Empirically, these optimization problems are much harder than just checking the existence of a feasible assignment of a given set of stations: in practice, we have been unable to exactly solve any value maximization problem at a national scale, even given weeks of time on high-performance computer clusters.

In theory, if we used a VCG mechanism that is based on an approximate optimization algorithm, that would make it nearly optimal for bidders to bid their true values, provided that the approximation is good enough. The problem with that approach is that, with the approximations that were achievable in practice, the "approximate VCG" price to a given winning station could differ by several orders of magnitude from its real VCG price. With such large pricing errors, truthful bidding could be far from optimal.

Finally, even if a sufficiently good approximation to VCG were possible, bidders might be unconvinced. In practice, bidders would be unable to independently verify whether their bids should win or lose or what their VCG prices should be. Also, even if they could compute well enough to do that, the law requires that bids be kept confidential, and therefore, the data for any verification would be unavailable. If bidders are uncertain about whether the FCC can or will compute sufficiently well, they could be inclined to manipulate their bids. The success of the auction hinged on the participation of small broadcasters who were not familiar with hard computational problems or auction mechanisms in general. If such broadcasters found the auction's computations too confusing, they might decline to participate at all.

A Heuristic Clock Auction. When high efficiency is an important goal and optimization is impossible, what is a market designer to do? Happily, for the incentive auction, we found a heuristic that achieves a high level of efficiency and meets all of the demands described above: it is computationally tractable and reasonably efficient, and the associated auction is convincingly ("obviously") strategy-proof. We can show this by a mix of theoretical analyses and the empirical analysis of extensive computational simulations.

The heuristic and auction that the FCC adopted belong to the class of "deferred acceptance" algorithms. In this auction, the auctioneer solicits offers from parties and iteratively rejects individual offers, possibly giving the rejected party an opportunity to submit another improved offer. When the auction ends, all offers that have not explicitly been rejected are accepted.

The design separates bidding into two interlocking components: the "reverse auction," in which the government buys television broadcast rights, and the "forward auction," in which the

government sells wireless licenses created from the former television broadcast frequencies.

The reverse auction is a descending clock auction, which means that the auctioneer quotes a declining sequence of prices to television stations, like a timer ticking down. From a bidder's perspective, the first step is to decide whether to participate at the (high) opening price. Then, each time its price is reduced, the station must either indicate that it would be willing to sell at that price or exit the auction irreversibly. When the auction ends, each participating station that has not exited sells its broadcast rights at the last price that it had agreed to offer.

Although a bidder may not understand all of the computations in the auction, each bidder knows enough to bid optimally. It knows that exit is irreversible, that its price can only be reduced (never increased), and that, after any price reduction, it will have another opportunity to exit. Then, regardless of how the clock prices are computed and how others bid, a bidder's optimal strategy is to exit only when the clock price falls below its value. Such a descending clock auction is said to be "obviously strategy-proof" (13).[†] In contrast, the optimality of truthful bidding in a VCG auction depends on very precise computations that bidders cannot verify and the bidder believing that no other bidder can peek at its bid. A bidder that cannot verify the auctioneer's computations or the privacy of its bids cannot be certain that truthful bidding in a VCG auction is optimal.

Feasibility Checking. To guarantee that the auction outcome is feasible, the reverse auction algorithm must never reduce a station's clock price if the station's decision to exit would cause the auction outcome to be infeasible. To verify that the outcome is feasible, the auction must identify a way to assign channels to all of the stations that either have exited or did not bid in the auction without violating any interference constraint. The FCC auction performs this verification using a "feasibility checker," which is software that determines whether a specified set of stations can be assigned to a given set of channels in a way that satisfies the millions of interference constraints.

We investigated many out of the box solutions for feasibility checking, notably two prominent commercial mixed integer programming (MIP) packages (CPLEX and Gurobi) and 19 open source solvers for the Boolean satisfiability (SAT) problem included in the ACLib library (15). Generally, we found that the SAT solvers outperformed the MIP packages; however, none performed nearly well enough to make us confident about deploying it in a real auction. We thus worked to design a custom solution based on existing SAT solvers. A key insight was that we did not need good worst case performance for SAT solving or indeed, graph coloring encoded as SAT; we needed a solver that performed well only on the sorts of problems that would arise in the incentive auction. All of these would involve subsets of the same 2,993 stations, among which the same pairwise interference constraints would hold. Furthermore, although we did not know stations' true values, we had a credible probabilistic model of these valuations in which values grow roughly linearly with population served and adjusted for population, correlate with the values of other stations in the same geographic area. Assuming this structure for station valuations could lead to regularities in the

sorts of feasibility checking problems encountered in the auction. For example, we might see stations serving densely populated cities exiting the auction before stations serving rural areas. More than a million feasibility checking problems generated by more than a hundred auction simulations provided data that reflected the regularities associated with the common constraints and the pattern of values.

A critical challenge was to build a feasibility checker optimized to perform well on this enormous dataset. Our approach used human intuition only to decide which design approaches could be considered, taking advantage of our domain-specific knowledge. Then, we used an automated search process for "algorithm configuration" (16) that leverages methods from machine learning and optimization along with a large computer cluster and weeks of compute time. The software searched our gigantic space of candidate algorithm configurations to identify one that worked well on our dataset.[‡] The result of the search was an "algorithm portfolio," dubbed SATFC (for SAT-based feasibility checker), consisting of eight SAT solvers that were constructed to complement each other well when run in parallel and that incorporate several different domain-specific algorithm ideas (19).

We also exploit parallel computation across runs of the feasibility checker. To explain how it works, we must first describe the timing and sequential processing of the reverse auction in more detail. After each tick of the clock in the auction, there is a period during which bids may be submitted followed by a period for bid processing. Broadcasters' bids are "processed" sequentially in some predetermined order. When station s 's bid is to be processed, SATFC attempts to determine whether s can feasibly exit. If the answer is negative or if the attempt times out, then s 's price is "frozen," and its bid is not examined. Otherwise, the software refers to s 's bid to see whether s has bid to exit in response to a lowered price. If so, s exits the auction; otherwise, s 's price is lowered, and the set of stations to continue broadcasting remains unchanged.

During the auction, any station s exits at most once. Many times, when a station is processed, it does not exit, and therefore, the set S of stations to continue broadcasting remains unchanged. If the next k stations to be processed are s_1, \dots, s_k , then k instances of SATFC can check whether each can be feasibly added to the set S . Let j be the index of the first station to exit. All of the first j parallel checks are usable by the auction, substantially reducing the clock time needed for processing. When S is augmented to include s_j , the parallel computations for the remaining $k-j$ stations become obsolete, and new checks are begun.

Handling Timeouts. To ensure that the reverse auction concludes within a reasonable amount of time with relatively small price decrements, the FCC decided that the clocks would need to tick down at least twice per 8-h business day. An analysis of auction simulations indicated that each feasibility check could, therefore, be given a maximum of 1 min to solve each problem. We evaluated SATFC's performance on data from new simulations and found that it could solve over 96% of these problems within a 1-min time limit compared with roughly 80% for the best off the shelf SAT solver and less than 40% for the best MIP package (Fig. 2). This reduction in so-called timeouts is dramatic but still

[†]This strategy proofness property holds only for broadcasters who have a single station for sale. For an owner of multiple stations, withholding one station from the auction could possibly raise the prices received by other stations (14).

[‡]Our overall approach is therefore similar to the way that recent "deep learning" algorithms construct classifiers that can label a dataset of images with high levels of accuracy (17). We thus call our approach "deep optimization" (18).

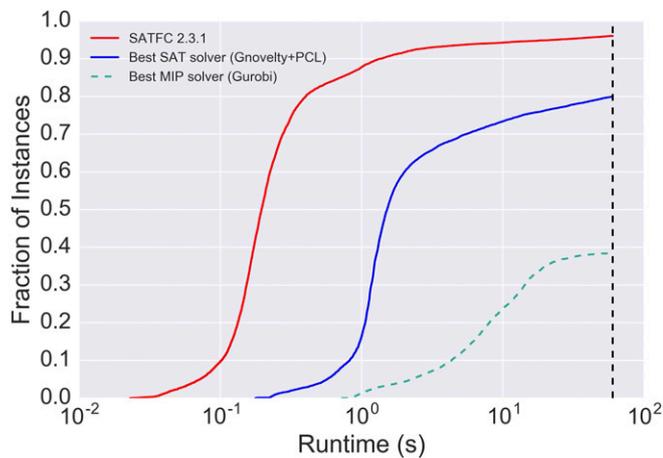


Fig. 2. Performance of the best MIP solver (Gurobi), the best off the shelf SAT solver (Gnovelty+PCL), and SATFC 2.3.1. The graph shows the fraction of problems solved within amounts of time ranging from a hundredth of a second to a minute.

not perfect—as, indeed, we expected given the computational hardness of repacking.

It is thus critically important that the heuristic clock auction deals gracefully with occasional timeouts (recall that VCG, for example, does not). In the clock auction, the solution is simple: when the feasibility checker is unable to establish the feasibility or infeasibility of a repacking problem, the auction simply treats that problem as infeasible and freezes the corresponding station's price. That solution slightly increases the total cost of buying stations for each unsolved feasible problem but does not affect the auction's feasibility or obvious strategy-proofness.⁵

Although it is necessary to set a time limit for each phase of the reverse auction, setting a maximum timeout for SATFC's unsolved problems has a cost: the bid processing period can end with both unsolved problems and considerable time remaining. Uncertainty about SATFC run times makes it impossible to avoid this difficulty if stations must be processed in a predetermined order. We now describe an elegant solution to this problem (which sadly, came after the auction design was finalized and hence, was not used in the incentive auction). The key idea is to determine the order in which bidders are processed dynamically. We begin by simultaneously checking the feasibility of assigning a channel to each station as if each was the first station in the ordering. As soon as any one checking problem completes, say for station s , that station is chosen to be the first station in the ordering. If s exits, then the set of stations S to be assigned channels is augmented to include s , and the feasibility checking problems for all remaining stations are restarted; otherwise, the other runs continue unchanged. We repeat this process until we reach the time limit for bid processing. The result is that no time available for feasibility checking is wasted. Either all stations are successfully checked or all easy stations are checked and the whole of the remaining time budget is spent simultaneously considering all of the hard, unsolved problems. This method does not avoid the computation time penalty that must be paid when a station exits the auction, and feasibility checks must be restarted for all of the still

⁵In addition, SATFC was designed to perform faster on feasible problems than on infeasible ones, and therefore, most of the timeouts occur on infeasible problems, for which the cost of a timeout is zero.

unprocessed bids. However, dynamic ordering ensures that, when we encounter such a station, no other station could be processed more quickly, and therefore, we pay the smallest possible penalty before restarting.

Setting the Clock Prices. What price should we quote to each station each time the clock ticks? If all stations that can feasibly exit were quoted the same clock price, the result would be that the stations demanding the highest compensation would exit first, except for the stations that can no longer feasibly exit. In some special cases, this “greedy heuristic” for choosing which stations continue to broadcast would yield an efficient outcome. For example, if all stations had the same location and broadcast power, then there would be some number K , such that any combination of K stations could feasibly be repacked. In this case, the auction would choose the K highest value stations to continue broadcasting while buying other stations' broadcast rights. In general, this clock auction algorithm results in the efficient outcome for all possible station values if and only if the sets of stations that can be feasibly repacked form a mathematical structure called a matroid (20), which is also equivalent to the stations being substitutes (21).

The all stations substitutes condition is not precisely satisfied in the FCC problem, and in simulations, the simple greedy heuristic is not typically optimal. In simulations, efficiency is enhanced compared with the simple heuristic by packing more stations with fewer interference constraints. The auction design can promote such outcomes by setting lower clock prices for those stations, encouraging them to exit earlier in the auction. Also, the expected total price paid to clear a fixed number of channels tends to be reduced by offering lower prices to stations serving fewer viewers, because those stations often have lower values for their broadcast rights and may be willing to sell for these lower prices.

The FCC auction design incorporates both of these ideas. It assigns to each station j a “volume” of $s_j = (P_j/l_j)^{1/2}$, where p_j is the population covered by j 's broadcast signal, and l_j is the number of stations with which j 's signal may interfere. A single “base clock price” descends during the auction, and the price quoted to each feasible station is computed by multiplying the base clock price by s_j . At some point in the auction, it may no longer be feasible for station j to continue broadcasting given the commitments that have already been made. At that point, j station's price is frozen: it declines no farther.

We illustrate the heuristic clock auction in our three-station example. Suppose that stations N and S both have volumes of one and station C has volume w . (Let us say that $w > 1$, because C has more interference constraints than N or S and perhaps, also covers a greater population.) In an auction round when the base clock price is b , stations N and S are quoted price b , and station C is quoted price wb . Suppose that b starts high and descends continuously. There are two cases to consider. If $v_C > w \max\{v_N, v_S\}$, then station C will be the first to exit, when the base clock price has fallen to v_C/w . Stations N and C become infeasible then, and therefore, there are no more price reductions. The auction stops, and the auctioneer acquires stations N and S , each of which is paid its clock price of v_C/w . If instead, $v_C < w \max\{v_N, v_S\}$, then whichever of N or S has the higher value exits first. Say that it is station N . When that exit happens, station C becomes infeasible, freezing its price at $w \max\{v_N, v_S\}$. Then, the base clock runs down—to zero if necessary—until station S exits. The efficient outcome is to acquire station C if and only if $v_C < v_N + v_S$, but the clock auction acquires C if and only if $v_C < w \max\{v_N, v_S\}$. Generally, no descending clock auction can guarantee an efficient

outcome when some sellers are complements, like stations *N* and *S* are in the example.

Luckily, although the actual US interference constraints make some stations complements, the tendency for stations to be substitutes is strong enough for the FCC clock auction to yield nearly efficient outcomes. In work with a collaborator (22), we conducted auction simulations using a small-enough set of stations for exact efficient channel assignments to be computed (stations within two interference links of Manhattan, which included all stations in Boston, Philadelphia, and Washington, DC). We defined the efficiency loss ratio of a mechanism *m* as the sum of values of stations winning and going off air under *m* divided by the sum of values of stations winning in the efficient (VCG) allocation. In our simulations, the FCC auction design had efficiency loss ratios between 0.8 and 10.0%, with costs (total payments to broadcasters) varying between 70 and 86% of the costs paid by VCG. The simulations also showed the importance of good feasibility checking: switching from SATFC to a simple “greedy” algorithm degraded the efficiency loss ratio by between 26 and 70% and increased costs by between 44 and 91%. Finally, setting the same clock price to all stations instead of setting prices proportional to FCC’s volumes would have raised the cost by between 36 and 54% (while having a small and ambiguous effect on efficiency). Follow-up work performed auction simulations at the national scale (18). While it was not possible to determine VCG allocations for these larger problems, this work showed that switching from SATFC to the greedy algorithm degraded the efficiency loss ratio by between 145 and 233% and increased costs by between 228 and 299%. These findings confirm the intuition that as the number of stations participating in the auction grows, it becomes increasingly important to use a strong feasibility checker.

Clearing Target. So far, our discussion of the reverse auction has presumed that we know how many channels to clear. In reality, the auction had to determine that as well, in such a way that the buyers pay at least enough to compensate the sellers, leaving no deficit. The auction’s approach to this is inspired by McAfee’s (11) design for homogeneous-good markets. Specifically, the incentive auction was designed to run through a series of “stages,” each with a different spectrum-clearing target. In the first stage, the clearing target was set to the highest nationwide target possible given the initial participation commitments by broadcasters at the opening prices, which happened to be 21 television channels (126 MHz of television spectrum), and which would have been sufficient to make 10 wireless bands available in most of 416 regions (100 MHz usable for wireless services plus some guard bands). The reverse auction in stage 1 closed at a total price of \$86.4 billion. Next, the forward auction began with low prices

for the wireless licenses in 476 regions. In this ascending auction, prices were increased for licenses for which there was excess demand until no excess demand remained. Then, because the \$23.5 billion of forward auction revenue from the first stage was too low to cover the cost of the reverse auction, the clearing target was reduced, and the forward and reverse auctions were continued with a smaller spectrum-clearing target. Reducing the clearing target provided additional channels for television broadcasting, making it possible to further reduce prices to many of the relatively expensive broadcasters. As a result, some of them exited, and others accepted lower prices. If television spectrum was a homogenous good, that would reduce the average price per acquired channel, and it tends to do the same in the actual auction. Similarly, the reduced target reduces the number of wireless licenses available in almost every area, allowing those prices to go up until demand falls to be equal to the reduced available supply. This process was designed to continue through several stages until the cost of the reverse auction is finally covered by the revenue from the forward auction (plus about \$2 billion for FCC expenses, including the estimated cost of postauction station repacking). It was theoretically possible for the auction to go through all of the possible stages and end with the number of channels to be cleared falling to zero (that is, no spectrum being reassigned).

Conclusion

The incentive auction’s closing conditions were met on January 18, 2017, with the auction being in stage 4, in which television channels 38–51 would be reassigned. The outgoing FCC chairman Tom Wheeler made the following statement:

The world’s first spectrum incentive auction has delivered on its ambitious promise. Reaching the Final Stage Rule means the benefits of the auction are indisputable. We will repurpose 70 MHz of high-value, completely clear low-band spectrum for mobile broadband on a nationwide basis. On top of that, 14 MHz of new unlicensed spectrum—the test bed for wireless innovation—will be available for consumer devices and new services. The auction will provide \$10.05 billion to broadcast television licensees who participated and billions toward deficit reduction.

More precisely, because wireless companies will pay in excess of \$19 billion total for their acquired spectrum, on subtracting the payments to broadcasters and the FCC’s expenses, the auction will net over \$7 billion for the US Treasury. While being unique in a number of ways, we believe that the auction offers a good example of how recent advances in economic theory and computer science can be combined to design radically new marketplaces, unlocking substantial economic value and benefiting all market participants as well as the US public as a whole.

1 Hayek FA (1945) The use of knowledge in society. *Am Econ Rev* 35:519–530.

2 Roth AE (1984) The evolution of the labor market for medical interns and residents: A case study in game theory. *J Polit Econ* 92:991–1016.

3 Milgrom PR (2004) *Putting Auction Theory to Work* (Cambridge Univ Press, Cambridge, UK).

4 Coase RH (1959) The federal communications commission. *J Law Econ* 2:1–40.

5 Vickrey W (1961) Counterspeculation, auctions, and competitive sealed tenders. *J Finance* 16:8–37.

6 Clarke EH (1971) Multipart pricing of public goods. *Publ Choice* 11:17–33.

7 Groves T (1973) Incentives in teams. *Econometrica* 41:617–31.

8 Myerson R (1982) Optimal coordination mechanisms in generalized principal-agent models. *J Math Econ* 11:67–81.

9 Myerson R, Satterthwaite M (1983) Efficient mechanisms for bilateral trading. *J Econ Theory* 29:265–281.

10 Mailath GJ, Postlewaite A (1990) Asymmetric information bargaining problems with many agents. *Rev Econ Stud* 57:351–367.

11 McAfee RP (1992) A dominant strategy double auction. *J Econ Theory* 56:434–450.

12 Segal I, Whinston M (2013) Property rights. *Handbook of Organizational Economics*, eds Gibbons R, Roberts J (Princeton Univ Press, Princeton), pp 100–158.

- 13 Li S (2015) Obviously strategy-proof mechanisms. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2560028. Accessed June 12, 2017.
- 14 Doraszelski U, Seim K, Sinkinson M, Wang WP (2016) Ownership concentration and strategic supply reduction. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2766468. Accessed June 12, 2017.
- 15 Hutter F, et al. (2014) AClib: A benchmark library for algorithm configuration. *Learning and Intelligent Optimization: 8th International Conference*, eds Pardalos PM, Resende MGC, Vogiatzis C, Walteros JL (Springer, Cham, Switzerland), pp 36–40.
- 16 Hutter F, Hoos HH, Leyton-Brown K (2011) Sequential model-based optimization for general algorithm configuration. *Learning and Intelligent Optimization: 5th International Conference*, ed Coello CAC (Springer, Berlin), pp 507–523.
- 17 LeCun Y, Bengio Y, Hinton G (2015) Deep learning. *Nature* 521:436–444.
- 18 Newman N, Fréchette A, Leyton-Brown K, Deep optimization for spectrum repacking. *Commun ACM*, in press.
- 19 Fréchette A, Newman N, Leyton-Brown K (2016) Solving the station repacking problem. *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence* (AAAI Press, Palo Alto, CA), pp 702–709.
- 20 Oxley J (2011) *Matroid Theory*, Oxford Graduate Texts in Mathematics (Oxford Univ Press, Oxford), Vol 21.
- 21 Milgrom PR (2017) *Discovering Prices: Auction Design in Markets with Complex Constraints* (Columbia Univ Press, New York).
- 22 Newman N, Leyton-Brown K, Milgrom P, Segal I, Assessing economic outcomes in simulated reverse clock auctions for radio spectrum. arXiv:1706.04324.