Incentive Auction Design Alternatives: A Simulation Study

A Preprint

Kevin Leyton-Brown
Department of Computer Science
University of British Columbia
kevinlb@cs.ubc.ca

Paul Milgrom
Department of Economics
Stanford University
milgrom@stanford.edu

Neil Newman
Department of Computer Science
University of British Columbia
newmanne@cs.ubc.ca

Ilya Segal
Department of Economics
Stanford University
ilya.segal@stanford.edu

March 8, 2022

Abstract

Over 13 months in 2016–17 the US Federal Communications Commission (FCC) conducted an "incentive auction" to repurpose radio spectrum from broadcast television to wireless internet. This paper revisits the descending clock "reverse" auction used to procure broadcast rights, taking a computational perspective. We describe an investigation into the quantitative significance of various aspects of the design based on extensive simulations, leveraging a reverse auction simulator and realistic models of bidder values.

1 Introduction

Over 13 months in 2016–17 the US Federal Communications Commission (FCC) conducted an "incentive auction" to repurpose radio spectrum from broadcast television to wireless internet. The result of the auction was to remove 14 UHF-TV channels from broadcast use, to sell 70 MHz of wireless internet licenses for $19.8 billion, and to make 14 MHz of spectrum available for unlicensed uses. With fewer remaining channels for TV stations, the TV spectrum needed to be reorganized; stations interfere with each other, so not all of them could be reassigned channels in the compressed TV band. Each station was either "repacked" in the leftover channels or voluntarily sold its broadcast rights, either going off the air or switching to a lower-quality band. These volunteers received a total of $10.05 billion to make repacking possible by yielding or exchanging their rights.

This paper uses a computational lens to revisit a key element of the incentive auction design: the descending clock "reverse" auction used to procure broadcast rights. We investigated the quantitative significance of various aspects of the design by running extensive simulations, leveraging a reverse auction simulator and realistic models of bidder values.

What is the value in taking another look at the incentive auction? Because the design was both novel and extremely complex (Leyton-Brown et al. 2017), it was not possible to thoroughly consider every potential design variation before the auction was run. Our goal in this paper is to gain new insights into how well the auction design performed, particularly asking which elements of the design were most important and which variations of the design might have led to even better outcomes. More broadly, we hope that our work will serve as an example for how simulations can be employed to understand and evaluate alternative market designs in complex settings.

The auction mechanics will be described in detail in Section 2, but roughly, the reverse auction worked by approaching stations in a round-robin fashion with a sequence of decreasing price offers for their broadcast rights. When a station refused an offer, it "exited" the auction irrevocably and was guaranteed a channel to continue broadcasting on after

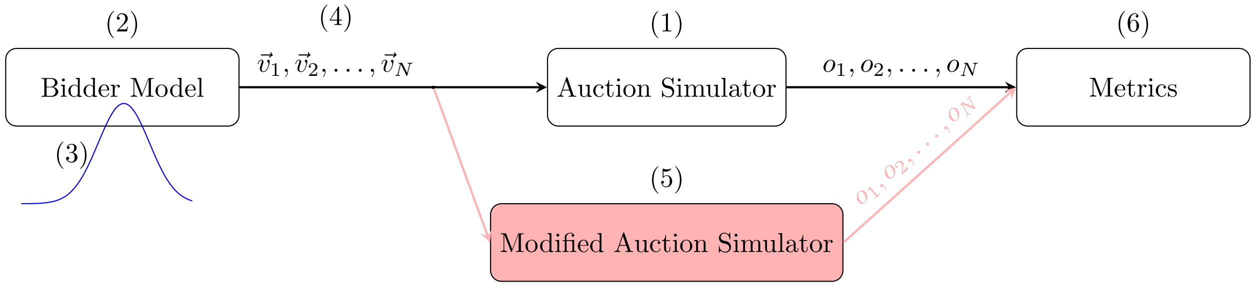

Figure 1: A high level schematic of our general simulation methodology. Labels correspond to the step descriptions. The same set of N samples from the bidder model are fed into both the auction simulator and a modified auction simulator. Each simulation produces an outcome o. The two sets of raw auction outcomes are then converted into metrics and compared on a per-sample basis.

the auction. As prices fell and more stations declined offers, the remaining channels filled up. Before processing any station's bid, a "feasibility checker" ensured that the station could still fit in one of the remaining channels alongside all stations that had already exited without causing undue interference; if it could not, that station "froze". Frozen stations were not asked for further bids and their prices were not lowered. Stations that remained frozen when the auction concluded were paid according to the last offer they had accepted. The auction included a provision to allow stations to exchange their broadcast rights for cash plus a channel in a less desirable (VHF) band and also a procedure to determine how many channels to repurpose, or "clear" (referred to as the auction's *clearing target*) in case the cost of clearing many channels proved to be too high. Licenses in the cleared spectrum were sold to mobile carriers using a "forward" ascending clock auction.

Our retrospective analysis is based on computational simulations. The reverse auction design was strongly informed by theoretical analysis of simplified models; see a survey in Section 3. Simulations are well suited to checking whether such theoretical findings carry over to more complex, real-world settings, as well as to pinpoint where further improvements would be valuable. The FCC used simulations in the design process to verify robustness of many design elements, some of which were publicly described (FCC 2015b). Simulations are particularly well suited to assessing designs with little or no historical data; the incentive auction was the first of its kind and so, even after its conclusion we only have a single data point.

Our general simulation methodology involves a number of steps: (see Figure 1):

1. Build an auction simulator (choosing an appropriate level of abstraction, as auction rules are often incredibly complex)
2. Create a bidder model, exposing parameters that control both valuations and behavior
3. Establish a probability distribution over parameters of the bidder model
4. Identify many plausible auction scenarios by sampling repeatedly from this probability distribution
5. Run paired simulations by holding this population of scenarios fixed and varying one or more elements of the auction design
6. Compare outcomes across paired samples using predetermined metrics

We built the simulator used in this paper, and have made the code freely available online at `https://github.com/newmanne/SATFC/`. We investigated two bidder models: (1) the only fully specified model from the literature of which we are aware; and (2) a new model that we contribute ourselves based on publicly released bid data. We contrast the results obtained from the two models as a robustness check. For the comparison step, we focused on two key metrics for settings in which simulations all clear the same amount of spectrum: the total value of the stations removed from the airwaves (value loss) and the cost paid to stations. When assessing design elements that affected the amount of spectrum cleared, we assumed that it is preferable to clear more spectrum, following the public statements made by the FCC about the auction's intended goals.

Our analysis considers four categories of questions: three concerning economic design and one concerning algorithmic design.

1. *How much value was added by including the VHF option for broadcasters?* Adding this option substantially increased the complexity of the auction and undermined some of its desirable theoretical properties, like

obvious strategyproofness and group strategyproofness, but could have reduced the number and value of stations taken off the air and the cost of buying those stations' broadcast rights. In Section 6.1, we show that under straightforward bidding, not repacking the VHF band would have increased the auction cost by roughly 20–30% and also decreased efficiency by roughly 5–10% (with variation depending both on random sampling and the choice of value model).

2. *How was the auction's performance affected by the auctioneer's procedure for deciding the number of channels to clear, compared to what it might have been if the number of channels to be cleared was predetermined?* The actual clearing mechanism was novel and received little comment from participants. In Section 6.2 we show that the clearing procedure led to higher costs and less efficient outcomes than predetermining the amount of spectrum to clear. Across all of our experiments, the clearing procedure raised average value loss by 5–26% and average cost by 4–50% (with this variation depending on random sampling, the choice of value model used, and whether the VHF band was repacked). In Section 6.2.2 we show that a simple alternative clearing procedure could have performed better.

3. *How much were costs reduced by offering lower prices to stations that reached smaller populations of viewers instead of just offering the same price to all and letting head-to-head competition set the prices?* The reduced price offers, called "pops scoring", were politically contentious and condemned by some opponents as price discrimination. In Section 6.3, we show that our conclusions differ based on the value model. Under our new value model, pops scoring led to average costs 5% (2%) higher than head-to-head pricing when the VHF band was (was not) repacked. Under the value model from the literature, head-to-head pricing led to average costs that were 39% (5%) higher relative to pops scoring when the VHF was (was not) repacked.

4. *Was auction performance significantly improved by the FCC's use of a customized feasibility checker to determine whether a station could be repacked alongside the set of stations continuing over-the-air broadcasting? How large might that effect have been?* This question is important because the design of customized feasibility checkers required a nontrivial effort; such efforts should only be made in the future if they yield gains. We answer this question affirmatively in Section 6.4. We search over off-the-shelf alternatives and show that substituting the custom feasibility checker with the best alternative could have increased both average costs and value loss by more than 20%.

The rest of the paper proceeds as follows. Section 2 explains the reverse auction in detail. Section 3 surveys related work. Section 4 defines our our valuation model (both the valuation model from the literature and our new model based on real bid data), our bidding model, and the probability distributions over the parameters of each from which we sampled. Lastly, Section 6 explores each of the above questions in detail using simulations.

2 Reverse Auction Basics

In this section we describe the reverse auction and show where it fits within the context of the incentive auction. We begin by introducing the station repacking problem. Solving repacking problems is a key subroutine within the reverse auction loop. We then proceed to the auction rules.

2.1 Station Repacking

Prior to the auction, each television station $s \in S$ in the US and Canada[1] was assigned to a channel $c_s \in \mathcal{C} \subseteq N$. The set of channels $\mathcal{C}$ can be partitioned into three equivalence classes, referred to as *bands*. Listed in decreasing order of desirability, these bands are: UHF (channels 14–51), high VHF ("HVHF", channels 7–13) and low VHF ("LVHF", channels 1–6). We use $\text{pre}(s)$ to refer to a station's pre-auction band, sometimes called a station's home band.

Each station is only eligible to be assigned a channel on a subset of $\mathcal{C}$, given by a *domain* function $\mathcal{D} : \mathcal{S} \to 2^{\mathcal{C}}$ that maps from stations to these sets. The FCC determined pairs of channel assignments that would cause harmful interference based on a complex, grid-based physical simulation ("OET-69" FCC (2013)); this pairwise constraint data is publicly available (FCC 2015c). Let $\mathcal{I} \subseteq (\mathcal{S} \times \mathcal{C})^2$ denote a set of *forbidden station–channel pairs* $\{(s, c), (s', c')\}$, each representing the proposition that stations s and s' may not concurrently be assigned to channels c and c', respectively.

The goal of the incentive auction was to remove some broadcasters from the airwaves and assign the remaining stations new channels from a reduced set $\overline{\mathcal{C}} = \{c \in \mathcal{C} \mid c < \overline{c}\}$. This reduced set is defined by $\overline{c} \in \mathcal{C}$; each choice of $\overline{c}$ corresponds to some clearing target. The actual incentive auction ended with $\overline{c} = 37$, allowing the higher numbered channels to be used for other purposes.

[1]Canadian stations did not bid in the auction but could be reassigned new channels.

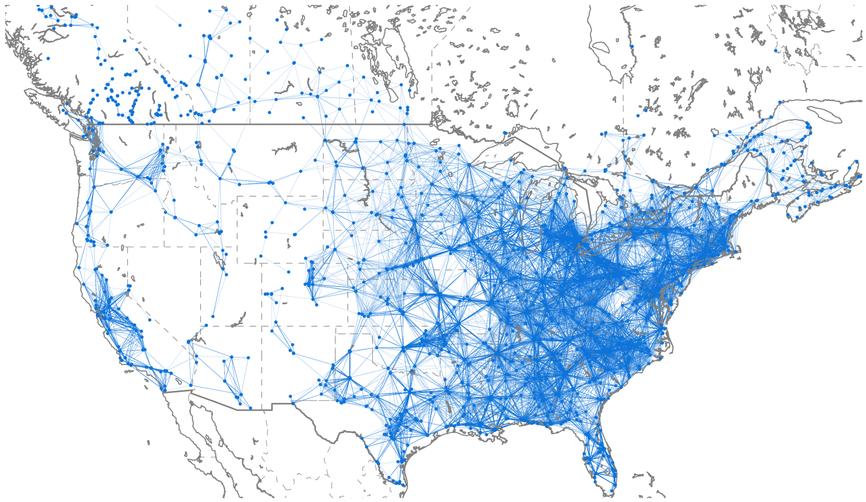

Figure 2: Interference graph derived from the FCC's constraint data.

A *feasible assignment* is a mapping $\gamma : \mathcal{S} \rightarrow \overline{\mathcal{C}}$ that assigns each station a channel from its domain that satisfies the interference constraints: i.e., for which $\gamma(s) \in \mathcal{D}(s)$ for all $s \in S$, and $\gamma(s) = c$ implies that $\gamma(s') \neq c'$ for all $\{(s, c), (s', c')\} \in \mathcal{I}$. As it turns out, interference constraints come in two kinds. *Co-channel constraints* specify that two stations may not be assigned to the same channel; *adjacent-channel constraints* specify that two stations may not be assigned to some pair of nearby channels. Thus, forbidden station–channel pairs are always of the form $\{(s, c), (s', c + i)\}$ for some stations $s, s' \in \mathcal{S}$, channel $c \in \mathcal{C}$, and $i \in \{0, 1, 2\}$.

Lastly, we define the *interference graph* as an undirected graph in which there is one vertex per station and an edge exists between two vertices s and s' if the corresponding stations participate together in any interference constraint: i.e., if there exist $c, c' \in C$ such that $\{(s, c), (s', c')\} \in I$. Figure 2 shows the interference graph for the US and Canada.

2.2 Reverse Auction

We now describe a simplified version of the reverse auction in which only the UHF band is repacked. In such a setting, only UHF stations participate in the auction and the only possibly outcomes for each station are going off air or continuing to broadcast in UHF. The real auction also repacked two VHF bands, but the inclusion of these bands complicates the auction rules significantly; we will briefly sketch some of these rule modifications later. The complete set of auction rules was published by the FCC in a 230-page document (FCC 2015a).

We begin by describing the reverse auction at a high level before giving more details about various key elements. First, stations respond to opening prices and decide whether to participate in the auction. Next, a solver finds an initial feasible channel assignment for all non-participating stations to minimize the number of channels required for those broadcasters, setting an initial clearing target $\overline{c}$. The auction then attempts to buy broadcast rights as necessary so that all stations remaining on air can fit into the available channels. It proceeds over a series of rounds, which consist of: (1) decrementing the clock and determining new prices, (2) collecting bids from each bidding station and (3) processing bids. We provide pseudocode for the reverse auction as Algorithm 1.

A forward auction to sell the cleared spectrum to wireless carriers follows the reverse auction. If sufficient revenue is raised in the forward auction, the incentive auction terminates; otherwise, the reverse auction resumes with a lower clearing target.

We elaborate on each step of the reverse auction below.

2.2.1 Prices

Prior to the auction, the FCC used a *scoring rule* to assign a *score* (also sometimes referred to as a *volume*). to each station, which we denote by $\text{score}\,(s)$. The score was used to determine individualized opening prices and was a function of both the station's interference constraints and the population of viewers that a station reached before the auction, which we denote by $\text{Population}\,(s)$. We will have more to say about scoring rules in Section 6.3.

Algorithm 1 Multi-Stage UHF-Only Reverse Auction Without Impairment

Input: S_{exited} contains initial set of non-participating stations; we assume no impairments. S_{bidding} contains the set of participating stations. $\overline{c}$ is the initial clearing target. $\mathcal{C}, \mathcal{I}, \mathcal{D}$ are the set of channels, the interference constraints, and the station domains respectively. FC is the feasibility checker. cutoff is the amount of time given to FC. p_0 is the starting clock price.

$S_{\text{winners}} \leftarrow \{\}$ ▷ Set of provisionally winning stations, initially empty
$\mathcal{P}(s) \leftarrow 0 \; \forall s \in S_{\text{bidding}}$ ▷ Provisionally winning prices; initially all 0
$\overline{\mathcal{C}} \leftarrow \{c \in \mathcal{C} \mid c < \overline{c}\}$ ▷ Channels available for repacking
$t \leftarrow 1$

while The incentive auction has not terminated **do**
  // Start a new stage
  $p_t \leftarrow p_0$ ▷ Reset clock prices
  while $|S_{\text{bidding}}| > 0$ **do**
    $p_t \leftarrow p_t - \max(0.05 \cdot p_t, 0.01 \cdot p_0)$ ▷ Decrease clock price

    // Check if any stations frozen in earlier stages have "caught up". Can be skipped in the first stage.
    for $\{s | s \in S_{\text{catchup}} \wedge p_t \cdot \text{SCORE}(s) \leq \mathcal{P}(s)\}$ **do**
      $S_{\text{catchup}} \leftarrow S_{\text{catchup}} \setminus \{s\}$
      if REPACKABLE($\{s\} \cup S_{\text{exited}}$) **then**
        $S_{\text{bidding}} \leftarrow S_{\text{bidding}} \cup \{s\}$
      else
        $S_{\text{winners}} \leftarrow S_{\text{winners}} \cup \{s\}$

    // Update prices
    for $s \in S_{\text{bidding}}$ **do**
      $P_{s;t} \leftarrow p_t \cdot \text{SCORE}(s)$

    // Bid processing
    Collect bid $b_{s;t}$ from each station in S_{bidding}
    for $s \in S_{\text{bidding}}$ **do**
      if REPACKABLE($\{s\} \cup S_{\text{exited}}$) **then**
        if $b_{s;t}$ = Exit **then** ▷ s bid to exit
          $S_{\text{bidding}} \leftarrow S_{\text{bidding}} \setminus \{s\}$
          $S_{\text{exited}} \leftarrow S_{\text{bidding}} \cup \{s\}$
      else ▷ s freezes
        $S_{\text{winners}} \leftarrow S_{\text{winners}} \cup \{s\}$
        $S_{\text{bidding}} \leftarrow S_{\text{bidding}} \setminus \{s\}$
        $\mathcal{P}(s) \leftarrow P_{s;t}$
    $t \leftarrow t + 1$

  Conduct a forward auction for the cleared spectrum $\{c \in \mathcal{C} \mid c \geq \overline{c}\}$
  Reverse auction cost $\leftarrow \sum_{s \in S_{\text{winners}}} \mathcal{P}(s)$
  if Revenue raised in forward auction $\geq$ Reverse auction cost **then**
    Terminate the incentive auction
  else
    $\overline{c} \leftarrow$ next clearing target ▷ Selected according to band plan
    $\overline{\mathcal{C}} \leftarrow \{c \in \mathcal{C} \mid c < \overline{c}\}$ ▷ Update channels available for repacking
    $S_{\text{catchup}} \leftarrow \{s \in S_{\text{winners}} \mid \text{REPACKABLE}(\{s\} \cup S_{\text{exited}})\}$ ▷ Find stations that may unfreeze
    $S_{\text{winners}} \leftarrow S_{\text{winners}} \setminus S_{\text{catchup}}$
// Incentive auction terminated
Pay $\mathcal{P}(s)$ to $s \in S_{\text{winners}}$

function REPACKABLE(S)
  Ask FC to find an assignment γ such that $\gamma(s) \in \mathcal{D}(s) \cap \overline{\mathcal{C}} \; \forall \; s \in S$, and $\gamma(s) = c \Rightarrow \gamma(s') \neq c'$ for all $\{(s,c),(s',c')\} \in \mathcal{I}$
  if A feasible assignment is found within cutoff seconds **then**
    Return True
  else
    Return False

The reverse auction is a descending clock auction. The initial base clock price was $p_0 = \$900$ both in the incentive auction and in our simulations. At the start of each round, the base clock price is reduced to $p_t = p_{t-1} - d_t$, where $d_t = \max(.05p_{t-1}, .01p_0)$. A station's score is used to translate the base clock prices to actual individualized station prices; that is, prices in each round $P_{s;t}$ are computed as $P_{s;t} = \text{score}(s) \cdot p_t$. We use $P_{s;\text{Open}}$ to refer to the auction's opening prices.

A station is said to be "winning" if it ultimately goes off air or moves to a different band at the auction's conclusion. More specifically, if the final channel assignment is γ, s is winning if $\text{post}(\gamma, s) \neq \text{pre}(s)$, where $\text{post}(\gamma, s)$ returns either the band to which s is assigned under γ or OFF if s is not assigned to a band under γ. We refer to the set of winning stations as S_{winners}. Throughout the auction, we track each station's "winning price" $\mathcal{P} : \mathcal{S} \rightarrow \mathbb{R}^+$. $\mathcal{P}(s)$ is the price that would have to be paid to s if the auction were to end immediately and s was a winning station. Initially $\mathcal{P}(s) = P_{s;\text{Open}}$.

2.2.2 Bidding

A station's bid in a given round corresponds to a binary decision when only the UHF band is being repacked. A station can accept $P_{s;t}$, indicating that it prefers to relinquish its broadcast rights and receive $P_{s;t}$. Alternatively, a station can reject $P_{s;t}$, indicating that it prefers to continue to broadcast. If a station's bid to reject an offer is ever processed (as will be explained in the following section), it is said to have "exited" the auction. Such a station is never asked to bid again. An exited station receives no compensation and will continue to broadcast in its pre-auction band after the auction (albeit on a possibly different channel). We refer to the set of exited stations by S_{exited}.

2.2.3 Bid Processing

In the bid processing step, stations are considered one after another, in descending order of $\frac{\mathcal{P}(s) - P_{s;t}}{\text{score}(s)}$. When a station s is considered, first the feasibility checker is invoked to determine whether it is possible to repack s along with the exited stations: i.e., given a time limit, it tries to find a feasible assignment for $\{s\} \cup S_{\text{exited}}$. If the feasibility checker cannot repack s, its bid is not examined, and s is said to be "frozen". In this case, $\mathcal{P}(s)$ is not reduced and s will no longer be asked to bid. If the feasibility checker can repack s, then $\mathcal{P}(s)$ is reduced and its bid is examined. If s bid to accept $P_{s;t}$, $\mathcal{P}(s)$ is lowered to $P_{s;t}$ and s remains "active", meaning it will be asked to bid again next round. If s bids to reject $P_{s;t}$, s permanently exits the auction.

2.2.4 Transitioning Between Auction Stages

A reverse auction stage ends when all stations are either frozen or have exited. Following each reverse auction stage is a forward auction stage where mobile carriers bid on licenses in the cleared spectrum. If the forward auction generates enough revenue to cover the costs of the reverse auction (the payouts to the winning stations $\sum_{s \in S_{\text{winners}}} \mathcal{P}(s))$[2], the incentive auction terminates and each frozen station is paid $\mathcal{P}(s)$. An unsuccessful forward auction (one that does not raise sufficient revenue) triggers another stage of the reverse auction with a smaller clearing target.

The incentive auction thus determines the amount of spectrum to clear endogenously by iterating through stages of reverse and forward auctions that clear progressively less spectrum until a stage occurs in which the forward auction covers the costs of the reverse auction. When a new reverse auction stage begins, $\overline{c}$ is increased, expanding the set $\overline{\mathcal{C}}$. Given additional channels, some frozen stations may now be repackable; such stations are said to be in "catch-up" mode. At the beginning of a new reverse auction stage, the base clock p_t resets. A station in catch-up mode "unfreezes" and returns to active status if it can be repacked in the first round in which it would face a weakly lower price than the price at which it froze, $\mathcal{P}(s)$. Subsequent stages otherwise proceed like the initial stage.

2.3 Worked Example

Consider the following worked example, illustrated in Figure 3. Consider an auction setting with six stations A, B, C, D, E, F with valuations $V_A > V_B > V_C > V_D > V_E > V_F$. Let each station be identically scored (so starting prices are the same for all stations); thus, we can drop the station subscript when discussing prices, writing P_t instead of $P_{s;t}$. Assume that the feasibility checker is perfect and will always correctly find a feasible assignment when one exists. For convenience, let clock decrements be so small that we can model the clock as falling continuously. Let $\mathcal{C} = \{1, 2, 3\}$ and let $\mathcal{I}$ be structured so that all stations have co-channel constraints on every channel with

[2]Technically, the forward auction had to generate about \$2 billion more: it also had to cover FCC expenses and the estimated costs of station retuning. In our examples and experiments, we ignore these additional requirements and only focus on the payments to winning stations.

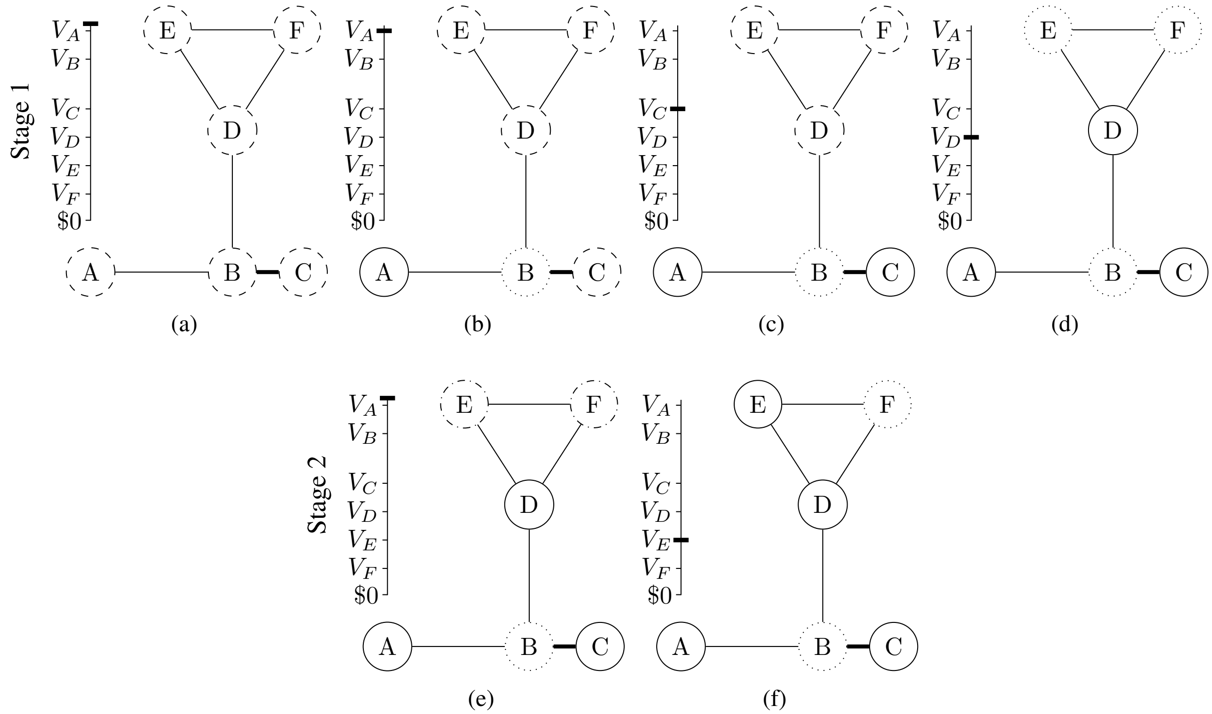

Figure 3: A worked example of a reverse auction involving six stations initially broadcasting on three channels. The stations are repacked into one channel in a first stage (top row) and then two channels in a second stage (bottom row). Times flows from left to right; the leftmost figure in a row shows the state at the beginning on the stage and a new figure is drawn each time a station exits. Each figure shows the interference graph, in which stations are nodes. Active stations are dashed, exited stations are solid, frozen stations are dotted, and stations with catch-up status are dashdotted. Connected stations have co-constraints, meaning that they cannot jointly broadcast on the same channel. B and C additionally have an adjacency-constraint, shown by the thick bold edge connecting them, meaning they cannot jointly broadcast on adjacent channels. To the left of each figure is the current clock price, with the values of each station marked.

each neighboring station according to the interference graph in Figure 3. Let stations B and C additionally have adjacency constraints with each other due to their close proximity, prohibiting them from jointly broadcasting on adjacent channels.

Let $\overline{c} = 2$ initially, so that the auction begins trying to repack the stations into a single channel. Near the high opening prices, stations bid to remain off air and their bids are processed because each station can exit. Nothing changes until $P_t < V_A$, at which point station A exits the auction. This movement freezes station B at price $\mathcal{P}(B) = V_A$, since A and B cannot both broadcast on the single channel. Other stations remain bidding and prices continue to fall until $P_t < V_C$ and C exits the auction. C's exit does not impact the other stations, so they continue to bid until $P_t < V_D$. At this point, D exits the auction, freezing E and F each $\mathcal{P}(E) = \mathcal{P}(F) = V_D$. Every station is now either frozen or has exited, so the stage ends. Stations B, E, and F are frozen at the end of the stage. If the incentive auction does not proceed to another stage, the total value removed from the airwaves will be $V_B + V_E + V_F$ and the total cost of freeing up two channels will be $V_A + 2V_D$.

Let us assume that in the forward auction, wireless carriers are unwilling to pay the repacking cost of $V_A + 2V_D$ to repurpose two channels worth of spectrum. The incentive auction then proceeds to a second stage.

In this second stage, let $\overline{c} = 3$, so only one channel is being cleared, with two channels available for repacking stations. Stations E and F can now be repacked alongside the exited stations and so enter catch-up mode. B cannot be repacked and remains frozen. P_t resets to a high value and then descends until $P_t < V_D$. At this point, E and F both transition from catch-up mode to bidding, since V_D was the price at which they froze. The auction continues until $P_t < V_E$, at which point E exits and freezes F at a price $\mathcal{P}(F) = V_E$. Again, all stations are either frozen or exited, so the second stage completes. Assuming the ensuing forward auction raises sufficient revenue, stations B and F will removed from the airwaves for a value loss of $V_B + V_F$ and a cost of $V_A + V_E$.

2.3.1 Extension: Repacking VHF

We conclude by noting some of the reverse auction modifications when the VHF band is also repacked. In multi-band auctions, stations can additionally bid to move between bands. Let $b_{s,t} \in B_{s,t}$ represent station s's bid in round t, where $B_{s,t} \subseteq \{\text{OFF, LVHF, HVHF, UHF}\}$ denotes the options available to s in round t. A station's options are restricted by a "ladder constraint", meaning that stations are never allowed to bid to move from "higher" bands (those occurring later in the list) to "lower" bands (those occurring earlier in the list). Additionally, stations are never allowed to bid for bands higher than their home band. In each round, rather than being offered a single price for going off the air, each bidding station is offered a separate price for each of its legal movements. We use $P_{s;b;t}$ to represent the price offered to station s in round t for selecting band b. A final difference we will mention is the introduction of fallback bids. While any number of stations can go off air, the VHF bands have limited capacity. As a result, the auction may not be able to accommodate certain bids for moving into VHF bands. For example, consider a UHF station s that bid for going off-air in the previous round and in the current round bids to move to HVHF. If s is not frozen when its bid is examined, the feasibility checker will try to determine if it is possible to fit s alongside the set of exited stations whose home band is HVHF. If the feasibility checker cannot find a feasible repacking, a fallback bid is used to determine what happens to s. This fallback bid can be either to exit or duplicate the previous bid (i.e., in the example, it would be as if s had again bid to go off air). We denote such bids as $\text{fallback}_{s,t}$.

3 Related Work

There exists an extensive literature on data-driven analysis and optimization of auctions; for example, on using past bids to learn an auction (Morgenstern and Roughgarden 2015) or to set personalized reserve prices (Golrezaei et al. 2017, Paes Leme et al. 2016, Roughgarden and Wang 2019, Derakhshan et al. 2019, 2020). To date there has only been a single incentive auction, and therefore only one sample of past bids for this auction design. Fortunately, simulations are particularly appropriate for examining the design of such markets that lack rich historical data. While there is relatively little literature on large-scale statistical analysis of the simulated behavior of candidate market designs in highly complex settings, of course this approach is widely applicable beyond the incentive auction. The international Trading Agent competition series (Wellman et al. 2007) used simulations that pit together researcher-contributed bots to gain general insights into autonomous bidding in electronic marketplaces. Another prominent recent example is the New South Wales fisheries combinatorial exchange, where simulations were used to determine if imposing linear and anonymous prices would result in acceptable efficiency losses (Bichler et al. 2019).

The incentive auction's design was strongly informed by theoretical analysis (Milgrom and Segal 2020); e.g., we know that when stations' bids are binary responses to a series of descending offers and when stations are all independently owned, the auction design is obviously strategyproof (Li 2017) and weakly group strategy-proof. Furthermore, when stations are all substitutes from the perspective of interference constraints, some descending clock auction can repack the efficient set of stations and another can implement the Myerson "optimal auction" (Milgrom and Segal 2020); however, here the substitutes condition does not hold exactly (the interference constraints are not a matroid). Theory on deferred acceptance auctions has continued to evolve since the auction: for example, (Gkatzelis et al. 2017) generalize deferred acceptance auctions to non-binary settings.

There were many questions about the incentive auction's design that could not be addressed via theory. As a result, simulations were used throughout every stage of the design process both by the FCC and the broader research community. Before the auction mechanism was even finalized, Kearns and Dworkin (2014) used simulations to characterize the space of feasible repackings based only on the interference constraints, for example estimating the minimum number of broadcasters that would have to relinquish licenses in order for a given clearing target to be feasible. A bit later in the design process, Cramton et al. (2015) used simulations to lobby for design changes such as changing the scoring rule and removing Dynamic Reserve Pricing, an idea that was ultimately scrapped. After the auction concluded, Doraszelski et al. (2017) used simulations to estimate how profitable and how risky bidder collusion strategies might have been. They leveraged their findings to to argue that the auction would have benefited from a rule change to deal with supply reduction strategies where an owner of multiple stations would face restrictions on which sets of stations they could use to participate in the auction. Ausubel et al. (2017) performed a post-mortem analysis of the incentive auction in a similar vein to our work, but with a primary focus on the forward auction; we discuss one of their proposed amendments to the clearing algorithm in Section 6.2.2.

The simulators used in the above projects varied greatly based on what was known about the auction at the time. Without yet knowing the auction mechanism, Kearns and Dworkin (2014) used a bidder model that consisted entirely of determining which stations would participate, which they studied using independent coin flips for every station as well as more complex models in which stations belonging to network affiliates made correlated decisions. Cramton et al. (2015) used a (non-public) valuation model developed through "discussions with many broadcasters, taking into account revenue data, historical station sales prices, station affiliation information, total market revenue, and other factors". Doraszelski et al. (2017) constructed a novel value model for stations in the reverse auction and made it publicly available; we discuss this model in Section 4.1.1. We leverage this model in our own simulations, and additionally constructed a second model based on auction bid data that was made public in 2019, two years after the auction concluded.

Beyond the value model, existing work taking simulation-based approaches to the reverse auction differs in the degree of fidelity with which it models the auction rules. In part, this was unavoidable since the design evolved over time: for example, early papers could not leverage the feasibility checker used in the auction because it did not yet exist, and so typically relied on off-the-shelf SAT solvers (e.g., (Kearns and Dworkin 2014)). The full auction rules are quite complex: while we also take some liberties outlined in Section 6, we are unaware of other work that runs incentive auction simulations as complex or realistic as ours (e.g., that include repacking the VHF band). Lastly, simulations are computationally expensive to run and previous studies differed in their computational capabilities. For example, experiments performed by Doraszelski et al. (2017) were restricted to the regional scale (considering subsets of at most 8% of UHF stations) and furthermore skipped certain feasibility checks ("limited repacking") due to computational constraints. Our simulator extends the one presented by Newman et al. (2017) and is closer to the actual reverse auction design than all of the previously discussed simulators. (We make no comment about the fidelity of the FCC's own simulator, since details about it are not publicly available.)

We conclude this section by noting two further papers that studied how prices should be set in a clock auction. First, Nguyen and Sandholm (2014) considered how to set prices to minimize expected cost, using the reverse auction as a test setting. Their methods reduce prices until feasibility is violated and then use a final adjustment round to regain feasibility. Their results are therefore not directly comparable to the FCC's design which maintains feasibility throughout and never increases stations' prices.

Second, Bichler et al. (2020) investigated the allocative efficiency of deferred acceptance auctions under different scoring rules in the problem of Steiner tree construction. It found that the choice of scoring rule affected the efficiency of the auction tremendously and that, under some scoring rules, deferred acceptance auctions were competitive with other mechanisms based on well studied approximation algorithms.

4 Value and Bidding Models

This section presents two separate value models and a model of bidding behavior that underly our simulations. The first value model follows Doraszelski et al. (2017); the second is based on our own analysis of bid data released after the auction by the FCC. An advantage of considering two different value models is that we were able to compare simulation results under both settings to assess the robustness of our findings. We conclude the section by describing a model of how stations bid as a function of these values.

4.1 Value Models

Each station s has a value $v_{s,b}$ for broadcasting in each permissible band b. We normalize so that a station has no value for being off air, i.e., $v_{s,\mathrm{OFF}} = 0$. Both models only provide $v_{s,\mathrm{UHF}}$, that is, home band values for UHF stations. For the two VHF bands in the auction, lower and higher VHF, we model a UHF station's value for switching to the HVHF band as $\frac{2}{3} \cdot v_{s,\mathrm{UHF}} \cdot \mathcal{N}(1, 0.05)$ and similarly for the LVHF band as: $\frac{1}{3} \cdot v_{s,\mathrm{UHF}} \cdot \mathcal{N}(1, 0.05)$—i.e., roughly two thirds and one third of the station's UHF value with some multiplicative Gaussian noise. [3] We generated values for stations initially in a VHF band by computing a hypothetical UHF value and then applying the fractional reductions for VHF bands just described.

4.1.1 The MCS Value Model.

Doraszelski et al.'s (2017) valuation model treats a station's value as the maximum of its cash flow value as a business and its stick value,[4] both of which were estimated from various sources including transaction data of station sales,

[3]We ensure that $v_{s,\mathrm{UHF}} > v_{s,\mathrm{HVHF}} > v_{s,\mathrm{LVHF}}$ by resampling if necessary. We rounded all values to the nearest thousand dollars. In rare cases when the value model provides an extremely low sample for $v_{s,\mathrm{UHF}}$, we set $v_{s,\mathrm{UHF}} = \$3000$. We describe robustness experiments on alternate parameter choices in Appendix A; our qualitative findings remain the same.

[4]The stick value represents the value of the broadcast license and tower, independent of the business; it can be more appropriate than cash flow when valuing non-commercial stations.

advertising revenue, and station features. We refer to this value model as the MCS (Max of Cash flow and Stick value) model for the remainder of the paper.[5]

The MCS model does not provide values for stations in Hawaii, Puerto Rico, or the Virgin Islands. We therefore excluded all of these "non-mainland" stations from our analysis completely when using both value models to preserve the same interference graph across simulations. We note that there are few such stations and they reach relatively small populations, so this exclusion is unlikely to have impacted our qualitative findings.

The MCS model also does not provide values for 25 mainland UHF stations. For these stations, we set $v_{s,\text{UHF}}$ to be proportional to population. We sampled the constant of proportionality from a log-normal distribution, where the mean and variance were calculated from samples of $\frac{v_{s,\text{UHF}}}{\text{Population}(s)}$ from other stations in that station's Designated Market Area (DMA)[6] or nationally if there were no other stations in the DMA.

4.1.2 A Novel Value Model based on Bid Data.

Two years after the incentive auction concluded, the FCC released data about the auction bids. The data contains the selected band and price (and fallback band and price when applicable) for each bid that was processed. The data contains the offers and set of bands for which each station s accepted opening prices, which we call $\text{PermissibleStartBands}_s$[7]. We use this data to construct a "realistic" model for station valuations. While the released bid data is not sufficient to reveal station values, it does make it possible to infer bounds on values. In some cases these bounds are relatively tight, with upper and lower bounds separated by a single clock interval, but most of the time they are looser, in some cases, to an extent that does not allow us to improve on the trival bound.

We inferred bounds on each UHF station's home band value, $v_{s,\text{UHF}}$. We began with the trivial bounds $0 \leq v_{s,\text{UHF}} \leq \infty$. We tightened bounds by applying the following rules[8] to the released bids:

1. $\text{OFF} \in \text{PermissibleStartBands}_s \implies v_{s,\text{UHF}} \leq P_{s;\text{OFF};\text{Open}}$
2. $\text{OFF} \notin \text{PermissibleStartBands}_s$[9] $\implies v_{s,\text{UHF}} \geq P_{s;\text{OFF};\text{Open}}$
3. $s \in S_{\text{winners}} \land \text{post}(s) = \text{OFF} \implies v_{s,\text{UHF}} \leq \mathcal{P}(s)$ [10]
4. $b_{s,t} = \text{OFF} \lor \text{fallback}_{s,t} = \text{OFF} \implies v_{s,\text{UHF}} \leq P_{s;\text{OFF};t}$
5. $(b_{s,t} = \text{Exit} \lor \text{fallback}_{s,t} = \text{Exit}) \land \gamma_t(s) = \text{OFF} \implies v_{s,\text{UHF}} \geq P_{s;\text{OFF};t}$

In words, we inferred that a station that included (did not include) starting off air as a permissible option had a value less than (greater than) its opening price for starting off air. A station that froze while bidding for off air and became a winner had a home band value less than its compensation. Whenever a station bid to remain off air, including as a fallback bid when attempting to move between bands, we inferred that the station's value was less than its price for remaining off air. Similarly, whenever a station bid to drop out of the auction (including as a fallback bid) while off air, we inferred that the station's value was greater than its price for remaining off air.

We then used these bounds to fit a model. We assume a model where value is proportional (in expectation) to population, $v_{s,\text{UHF}} = \text{Population}(s) \cdot n_s$. Here n_s is some number in units of \$/pop, sampled from an unknown cumulative distribution function (CDF) N. Our upper and lower bounds on a given station's home band value can be translated into upper and lower bounds on a station's \$/pop by dividing by s's population. Let x_s and y_s represent lower and upper bounds respectively on s's observed \$/pop sample n_s, i.e., $x_s = \frac{v_{s,\text{UHF,Lower Bound}}}{\text{Population}(s)}$, $y_s = \frac{v_{s,\text{UHF,Upper Bound}}}{\text{Population}(s)}$. Our goal is a non-parametric maximum likelihood estimate of the distribution function N. Note that by definition $N(y_s) - N(x_s) = \Pr(x_s \leq n_s \leq y_s)$. Our goal was to maximize the product of these terms subject to constraints ensuring that N is a valid CDF. Let Z be a list containing all of the x_s and y_s in ascending order, such that Z_1 is the smallest element and $Z_{2|S|}$ the largest. Our problem was then as follows.

[5]Our implementation follows the value model described in a 2016 draft of the paper; the authors have subsequently made revisions (e.g., to the way that population is counted for low-power non-commercial stations).

[6]Nielsen divides the US into 210 geographic DMAs (Nielsen 2022).

[7]Prior to the auction, stations accepted or rejected prices for each of their eligible bands; the optimization that determined the initial starting assignment was then constrained to only place stations on accepted bands.

[8]It is possible to further tighten bounds based on stations' bids for their VHF options; however, for some stations, this leads to contradictory bounds. We therefore discarded any bidding information related to VHF, including all bids made after a station (possibly) moved to a VHF band.

[9]This includes both non-participating stations and those that participated conditional on starting in a VHF band.

[10]We exclude station 35630, which received \$0 for its license.

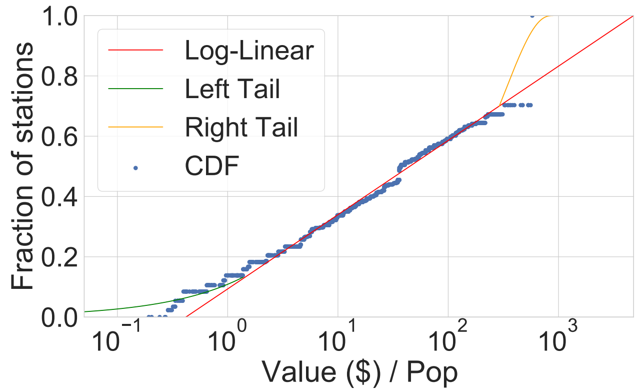

Figure 4: CDF values for our maximum likelihood estimate of N and a log-uniform distribution fit to the results. The left and right tails are Generalized Pareto Distributions.

$$\text{maximize} \prod_{s \in S} \big(N(y_s) - N(x_s)\big) \tag{1}$$

$$\text{subject to } N(x_s), N(y_s) \in [0,1] \, \forall s \in S \tag{2}$$

$$N(Z_1) \le N(Z_2) \le \cdots \le N(Z_{2|S|}) \tag{3}$$

Noting that taking the log of the objective function preserves its maximum, we minimized the negative log likelihood, $-\sum_{s\in S} \log\left(N(y_s) - N(x_s)\right)$. This is a constrained optimization problem with a nonlinear convex objective function and convex constraints. The constraints ensure that the values of N must fall between 0 and 1 and that N must be non-decreasing. The solution uniquely identifies the value of N at each point x_s and y_s. We translated this problem into a second-order cone program using `cvxpy` (Diamond and Boyd 2016) and solved it using the `ECOS` (Domahidi et al. 2013) solver. The results (Figure 4) suggest that N is a log-linear function. Data was sparsest in the tails of the distribution, especially in the right tail, so we replaced the log-linear segments in both tails with Generalized Pareto Distributions (GPDs). Specifically, we labelled the bottom 15% segment of the CDF the left tail and top 30% segment the right tail. The CDF of the GPD is $F(x) = 1 - (1 + \frac{x-\mu}{\sigma} \cdot \xi)^{\frac{-1}{\xi}}$ for $\xi \neq 0$, where μ, σ, ξ are the location, scale, and shape parameters respectively. We fit the left tail to the set of points Z falling in the region described (about 7% of our data). We did not have any data to fit the right tail of the distribution, so we fit a GPD using control points such that a $/pop value 12.5% higher than any we observed occurred at $y = 0.975$ and one 25% higher than any we ever observed occurred at $y = 0.99$. To generate UHF values, we multiply a station's population by a sample from the modeled distribution of N. (Concretely, to sample a station value from this model, we first sample u uniformly from (0,1), then convert it into a value $\exp\left(u \cdot (c - a) + a\right)$. We had $a = -1.1087, c = 8.6981$.) In what follows, we refer to this model as the BD (bid data) model.

While we recognize that the value distribution's right tail is the most *ad hoc* part of our model, we believe that the choices we made about this part of the distribution are unlikely to have substantially impacted our results. The reason is that stations with sufficiently high values are unlikely to choose to participate even at the incentive auction's opening prices, at which point it does not matter to the auction's outcome exactly how high those values are. Indeed, we calculated that only 3.4% of UHF stations nationally would have *any* chance of participating in the auction conditional on their values having been sampled from any part of the right tail. Since each station has a 30% chance of having its value sampled from the right tail in any given simulation, on expectation any change to this section of the curve would alter the bidding behaviour of fewer than 1% of UHF stations. Furthermore, it is likely that conditional on a station's value having been drawn from the right tail and the station having chosen to participate, these stations would exit early in the auction before interference constraints start to bind tightly, and that this behaviour would not depend on the exact shape of the right-tail distribution.

4.2 Bidding Model

We now describe our model of how stations bid in the auction.

A station participated in our simulations if its opening price for going off the air exceeded its value for continuing to broadcast in its home band, i.e., if $P_{s;\text{OFF;Open}} \geq v_{s,\text{pre}(s)}$. After excluding 64 non-mainland stations, we had 1813 stations eligible to participate in our simulations, composed of 1407 UHF stations, 367 HVHF stations, and 39 LVHF stations. We note that the incentive auction had 1030 participants (FCC 2019).

A key difference between the two value models is the participation rates that they elicit. In the MCS model, station values tend to be low relative to opening prices, leading to very high participation rates. For example, considering 10 000 sampled value profiles of UHF stations, 1196 of these stations participated in every sample, and only 46 had a less than 80% chance of participating. In contrast, when running the same experiment with the BD model, no station participated in every sample and average participation rates were 60%, closer to the 64% participation rate of UHF stations in the incentive auction.

In auctions with UHF options only, in which the only allowed bids are to remain off air or exit the auction, a single station faces a strategic situation in which its utility is "obviously" maximized by remaining off air if the price exceeds its value and by exiting otherwise—that is, by bidding myopically (Li 2017). The situation when VHF options are included is no longer obvious in this sense, but we continue to assume for simplicity that when bidding in round t, a station selects the offer that myopically maximizes its net profit, ${}_{b\in B_{s,t}} P_{s;b;t} + v_{s;b}$. When fallback bids are required, we again assume that stations select the option that maximizes their net profit. We note that in the released bid data, only 52% (13%) of bids to move into LVHF (HVHF) were successful. Given that stations seeking to move to a VHF band faced a meaningful risk that their bid might fail to execute, some bidders might have benefited by bidding on VHF bands before it was straightforwardly optimal to do so. However, despite the potential drawbacks of straightforward VHF bidding, we are not aware of any behavioral rule that can be applied to all bidders and that is arguably more realistic.

5 Simulator Design and Experimental Considerations

We now explain how we compare simulated outcomes and discuss some simplifications our simulations make relative to the real auction process. We conclude the section by describing some other elements of our experimental setup.

5.1 Metrics

One goal for the auction is efficiency: for any given clearing target, to maximize the total value of the stations that remain on the air, or equivalently, to minimize the total value of the stations removed from the airwaves. We focus on the latter definition—value lost instead of value preserved—because it is unaffected by value estimates for large, highly valuable stations that do not participate in the auction. That is, value preserved includes the values of easy-to-repack stations, even those that do not participate in any interference constraints, and leads to efficiency estimates near 100% when few stations go off air relative to the number that remain on air, even when the number of stations going off air is large relative to the number required by an optimal solution.

We define the *value loss* of an auction outcome as $\sum_{s\in\mathcal{S}} v_{s,\mathrm{pre}(s)} - v_{s,\mathrm{post}(\gamma,s)}$. Given an efficient repacking γ^* that minimizes the value loss for a given value profile, our metric for allocative efficiency is the value loss ratio $\frac{\sum_{s\in\mathcal{S}} v_{s,\mathrm{pre}(s)} - v_{s,\mathrm{post}(\gamma,s)}}{\sum_{s\in\mathcal{S}} v_{s,\mathrm{pre}(s)} - v_{s,\mathrm{post}(\gamma^*,s)}}$ where γ is the simulation's final assignment. In general, we do not have access to γ^* because it is too difficult to compute, so instead we resort to converting our absolute metric into a relative one and comparing the value loss between two simulations' final assignments (i.e., the ratio of value loss ratios, noticing that the denominators which depend on γ^* cancel out in this case).

Our second metric is the cost to clear spectrum: the prices paid to all winning stations, $\sum_{s\in S_{\text{winners}}} \mathcal{P}(s)$.

We use the terms "efficiency" and "cost" below as abbreviations that refer to value loss and the total payments made to broadcasters that go off air or change bands. Outcomes with low value loss and low cost are preferable. It is straightforward to compare two outcomes if they both clear the same amount of spectrum and one is both more efficient and cheaper; otherwise, any comparison requires a judgement call about how the two metrics should be traded off.

5.2 Impairments

The incentive auction's design requires it to begin from a feasible channel assignment for the non-participating stations. However, it may not always be possible to find a feasible repacking for the non-participating stations in the set of remaining chanels $\overline{\mathcal{C}}$ induced by the initial clearing target $\overline{c}$. The FCC's rules therefore allowed a small number of stations to be assigned to channels *within* the 600 MHz band, even though doing so degrades the desirability of the mobile broadband licenses sold in the forward auction. Such stations placed within the 600 MHz band are referred to as "impairing". There are two types of impairments: those caused by non-US stations that would be present even if every US station participated in the auction ("essential impairments") and those caused by non-participating US stations.

Despite our best efforts to make our simulations realistic, we could not replicate the optimization procedure the FCC used to determine the initial clearing target $\overline{c}$ and the set of impairing stations. The problem was that this optimization relied on data that is not publicly available: the Inter-Service Interference ("ISIX") constraints that determine which geographic areas are impaired when a station is placed on the 600 MHz band (i.e., the spectrum that will be resold, channels in $\mathcal{C} \setminus \overline{\mathcal{C}}$). Without the ISIX data, we also could not replicate the analogous optimizations that the FCC conducted between auction stages.

Instead of determining the inital clearing target via an optimization, except when otherwise noted, we started our simulations at the 84 MHz clearing target (the clearing target at which that the real incentive auction concluded) and ran auctions for only a single stage. We did this both because running multiple stages of the reverse auction is computationally expensive and because multi-stage simulations depend on additional assumptions about the forward auction. We do explore multi-stage auctions that begin from other clearing targets (including 126 MHz, the clearing target on which the incentive auction began) in Section 6.2.

Given a clearing target, our simulator determines which stations to impair as follows. Prior to running the auction, we raise the starting base clock price in 5% increments until we achieve full participation for UHF stations. We then solve an optimization problem to find an initial assignment γ that minimizes the number of essential impairing stations, breaking ties by minimizing the aggregate population of impairing stations. Stations marked as impairing do not interfere with each other or with stations assigned to channels in $\overline{\mathcal{C}}$. Having dealt with the essential impairments, we then run the auction starting from the inflated base clock price. While the base clock price is higher than its normal starting point, only UHF stations are asked to bid and the VHF band is considered "locked"—stations cannot bid to move into the VHF band. When the starting base clock price is reached, any station that would not have participated at the opening prices will have exited or be frozen. We consider any stations that have frozen at this point to be impairing, and so remove them from the problem entirely; we do not include impairing stations in any of our metrics unless explicitly noted. After the impairing phase, VHF stations make their participation decisions at the normal starting base clock price and the VHF band is unlocked. We repeat this optimization between each stage of the auction, so impairing stations can leverage the newly available spectrum from the reduced clearing target to possibly become non-impairing. Specifically, we use our optimization procedure to identify any removable essential impairments, then we proceed with the between-rounds transition as described in Section 2.2.4, except that p_t can be reset back to the impairment regime. Lastly, we note that there are no essential impairments given a clearing target of 84 MHz and that there are 9 at a clearing target of 128 MHz.

In the real auction, stations could choose to participate conditional on being initially assigned to one of a subset of bands. Another role of the initial clearing target optimization was to accommodate such stations. Since we could not replicate this optimization we did not allow stations to select their preferred starting bands. Instead, in our simulations, all participating stations start off air. In practice, 85% of stations indicated off air as their preferred starting band and only 5% of stations declined off air as a permissible starting band.

5.3 Some Other Experimental Considerations

As the reverse auction is simplest to reason about when only the UHF band is repacked, we ran both simulations that only repacked the UHF band and simulations that also repacked the VHF band. This allowed us to investigate which of our results generalize across settings. We ran simulations on both value models described in Section 4. Canadian stations were also involved in the repacking process (they could be moved to a new channel, but did not participate in the auction and could not be purchased). We included all 793 Canadian stations in our simulations. Unless otherwise stated, in every one of our experiments we took 50 samples per treatment. We gave feasibility checks 60 seconds to complete (as in the real auction, though of course we were unable to use exactly the same hardware) unless otherwise stated.[11]

6 Experiments

We are now ready to present our experimental results. As explained in the introduction, our experiments are divided into four categories, based on which element of the auction design they investigate: (1) repacking the VHF band in

[11] We note a subtlety regarding runtime measurements. Since the feasibility checker uses a walltime cutoff, there will be some degree of unavoidable noise (i.e., problems that require time very similar to the cutoff threshold). As a result, different auction trajectories starting from the same value profile can occasionally be caused by such measurement noise rather than a given design change being tested. This is difficult to control for, but the effect is random and averages out across samples.

addition to UHF; (2) choosing the order in which stations are processed via a scoring function; (3) determining a clearing target by iterating between reverse and forward auction stages; and (4) determining which stations to freeze by checking the feasibility of station repackings.

6.1 Repacking the VHF Band

The incentive auction reduced only the number of UHF channels, but repacked stations in three bands: UHF, HVHF, and LVHF. Repacking the two VHF bands offered the potential for cost savings and efficiency gains, as UHF stations might have been willing to accept a smaller payment to move to a VHF channel instead of going off the air and this could have constituted a net gain, even taking into account the possible need to compensate VHF stations for going off air to make space. An optimal repacking for a given value profile can only become weakly more efficient when the VHF bands are included as more configurations of stations become available.

However, adding extra bands to the reverse auction complicated an otherwise elegant design. To reduce strategic options available to bidders, a rule prohibited stations from moving from lower priced options to higher priced ones. This "ladder constraint", for example, prevented a station tentatively assigned to the HVHF band from bidding to go off air. The result was that stations no longer possessed obviously dominant strategies, as they could potentially benefit from choosing when to move up the ladder. Price calculations became more involved as each option had to be priced appropriately. Fallback bids were added to the bidding language to accommodate the case where a bid to move becomes invalid by the time it is processed. Also, unlike in UHF-only settings where freezing is permanent, VHF stations can freeze and later unfreeze within the same stage if other stations move out of their home bands, complicating bid processing.

All of this extra complexity made the auction more difficult to explain to station owners, which mattered since some of the participants were relatively unsophisticated and encouraging them to participate was a first-order concern. It is thus sensible to ask whether the additional complexity was worthwhile. While repacking VHF increased the space of possible assignments to potentially include more efficient allocations, it was possible that these improved allocations may only arise under exotic bidding behavior, or that they would have come at a high cost. What changes to efficiency and cost arise when VHF options are included and bidders bid straightforwardly?

To answer this question, we ran two sets of auctions: the first repacking both the VHF and UHF bands, the other repacking only the UHF band. Our results are shown in Figure 5. Each point represents the outcome of one simulation, with its x-axis position denoting its efficiency and its y-axis position denoting its cost. Since it is difficult to show graphically which auctions use the same paired value profiles for even a modest number of samples, rather than plotting raw efficiency and cost on each axis we instead plot normalized efficiency and cost. That is, we select one setting (in this case, auctions that repack VHF) as the reference treatment, and then plot the ratio of the efficiency (cost) of each simulation from additional treatments relative to the corresponding simulation using the same value profile in the reference treatment. With this choice, the reference treatment always corresponds to the point $(1, 1)$, represented in our figures by a diamond. For each treatment we also plot a star to indicate each metric's mean value.

Altogether the experiments took a little over one CPU year to run. Using both value models, we observed that repacking the VHF band led to a significant reduction in payments—on average, sampled UHF-only auctions cost 1.23 and 1.32 times as much as their VHF-repacking counterparts using the MCS and BD value models, respectively. The impact on efficiency was more modest and less uniformly positive. Under the MCS (BD) value model, sampled UHF-only auctions experienced 1.05 (1.12) times higher value loss on average. On the whole, our simulations suggest that if bidders continued to bid straightforwardly despite the complex design, repacking the VHF band was an important design choice that likely led to lower costs and also somewhat more efficient outcomes.

6.2 Multi-Stage Clearing

The incentive auction included both reverse and forward auctions in order to let market forces determine the appropriate amount of spectrum to clear. The FCC decided to run the reverse auction in each stage separately from the forward auction to economize on staff resources, and it chose to run the reverse auction first within each stage. Each successive stage cleared a smaller amount of spectrum for a lower overall cost than the previous stage. The actual incentive auction went through four stages, with four different clearing targets.

One obvious practical drawback of a multi-stage approach is its impact on auction length: running multiple stages takes time. As noted by Ausubel et al. (2017) "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".

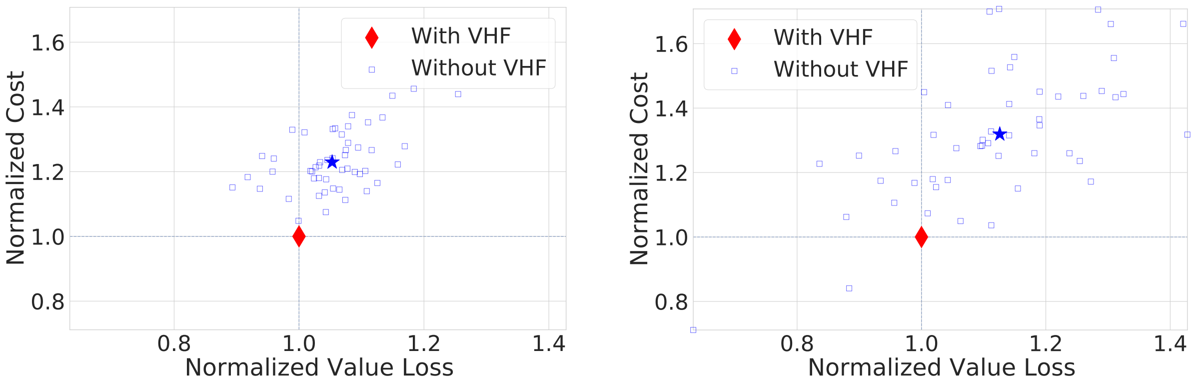

(a) MCS

(b) BD

Figure 5: Comparing auctions that only repack the UHF band against auctions that also repack VHF bands.

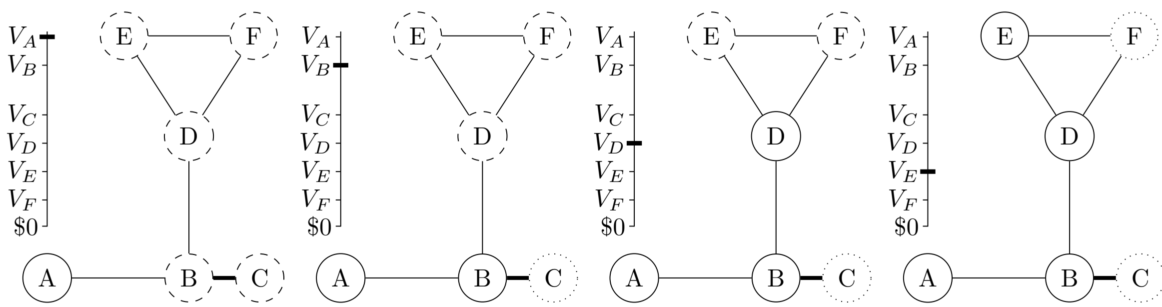

(a) (b) (c) (d)

Figure 6: Returning to the example of Figure 3, but now repacking two channels right from the outset.

A less obvious concern is how the multi-stage approach impacted the final outcome. While this multi-stage design made the amount of spectrum to clear endogenous, this property might have come at a cost, since irrevocable choices in the early stages were based on unrealizable clearing targets but could affect the final allocation.

For a concrete example of how such harm could occur, we return to the worked example in Figure 3 and ask what would have happened if the auction had initially set $\overline{c} = 3$ and tried to repack stations into two channels from the outset. We visualize this scenario in Figure 6. As before, A is the first to exit. However, B does not freeze this time around, because A and B are both repackable using two channels. When $P_t < V_B$, B will exit. At this point C freezes at $\mathcal{P}(C) = V_B$. D is the next to exit, followed by E, finally freezing F at $\mathcal{P}(F) = V_E$. Recall that the value loss and cost of the two stage auction were $V_B + V_F$ and $V_A + V_E$ respectively. In the single-stage alternative, the value loss is $V_C + V_F$ and the cost is $V_B + V_E$: the outcome is preferable according to both of our metrics!

In the example, the clock auction succeeds in optimizing the second stage target. In general, if the clock auction (nearly) optimizes for each clearing target and if the stations packed for the higher target are not a subset of those packed for the lower target, then one should expect a multi-stage auction to perform worse than an auction that starts by setting the correct target. The example suggests that the problem is that stations packed in late rounds of an early stage might cause problems in subsequent stages. Conversely, it is not always true that the auctioneer is better off predetermining the amount of spectrum to clear. In Appendix B we present a counterexample in which the multi-stage clearing procedure leads to a cheaper, more efficient outcome than would be achieved in a single stage with the correct clearing target.

We therefore turn to simulations, asking two questions: (1) what if any economic costs arose due to multi-stage clearing? (2) Would an "early-stopping" alternative to the reverse auction have performed better?

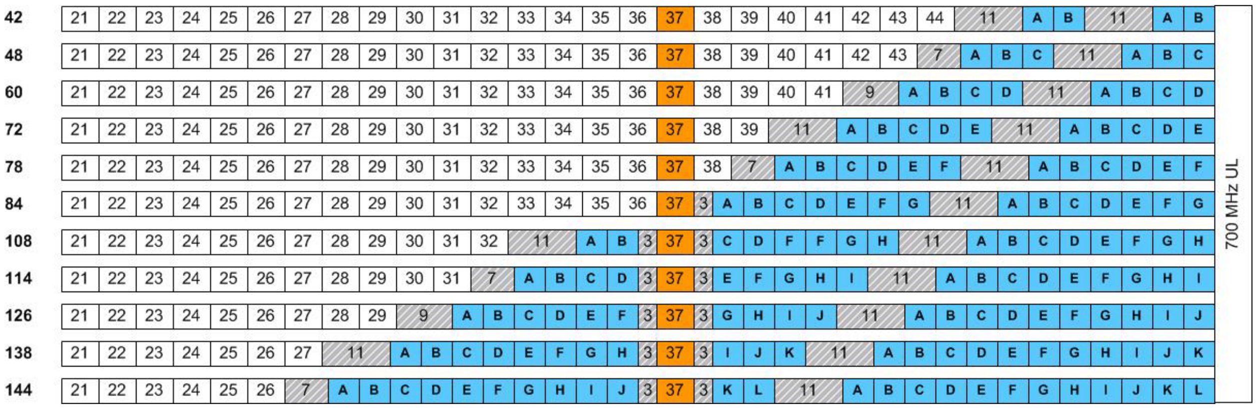

Figure 7: The FCC's band plan for 600 MHz. Each row corresponds to a different clearing target and shows which channels would be repurposed and which would remain for TV broadcasting.

6.2.1 The Economic Impact of Multi-Stage Clearing.

To assess the economic impact of multi-stage clearing, we ran experiments that begin trying to clear 126, 114, 108, and 84 MHz of spectrum, each proceeding to follow the ordering of clearing targets selected by the FCC (see Figure 7, leading to four-, three-, two- and single-stage auctions, respectively. These experiments took 31 years of CPU time.

A complicating factor in this comparison is that for any given value profile, the final set of impairing stations may differ based on the starting stage, which muddies the interpretation of cleared spectrum as a measure of performance. The impairment mechanism will attempt to make any impairing stations non-impairing in between each stage as more spectrum is added, but this does not guarantee that the cleared spectrum is comparable whenever impairing stations remain at the auction's termination. To improve the comparison, we analyze the results of these experiments in two ways. The first is to consider a world without impairments, corresponding to the case where the FCC is willing in principle to pay any amount to stations to eliminate impairments. We present such results in Figure 8. We observe that running the auction through multiple stages had downsides both in terms of cost and efficiency, especially when the VHF band was repacked. When the VHF band was repacked, in simulations using the MCS model, perfectly forecasting the clearing target might have roughly shaved a third off costs and only had 80% of the value loss of a four-stage solution. Under the BD model, four-stage simulations cost 1.19 times as much as their single-stage counterparts and had 1.14 times as much as value loss. The results are similar but less dramatic in magnitude when investigating UHF-only auctions. In UHF-only simulations using the MCS (BD) model, four-stage solutions cost roughly 5% (10%) more and had roughly 5% (10%) additional value loss compared to perfectly forecasting the clearing target. In all cases, even if the exact stage was not perfectly selected, starting the auction closer to the final stage would also have yielded significant improvements to both metrics on average.

A second way of analyzing the results is to look at the results including impairing stations. Of course, this leads to "apples to oranges" comparisons in the sense that the quality of cleared spectrum will vary across simulations. We do not have a reliable way of assessing how much a given impairment devalues spectrum. The answer surely depends on the location and power of the exact stations in question, but a blunt proxy number for how bad a set of impairing stations might be is the sum of their population (less is better). In fact, there were no impairing stations remaining at the final stage of any of the experiments using the MCS, so results remain the same for this model. Results for the BD model are visualized in Figure 9. Results have much higher variance than in the previous analysis. (Only) in the UHF-only case we no longer observe the trend that more stages led to worse outcomes on average—on average, four-stage simulations cost and had value loss about 1% higher than single-stage simulations, and two- and three-stage simulations performed a little better (1–2%) on average on each metric. When looking at the size of the impairing station populations, however, single-stage auctions resulted on average in the least impaired spectrum: four-stage, three-stage, and two-stage simulations had mean impairing populations of 206, 205, and 206 million respectively compared to 192 million for single-stage simulations. When the VHF band is repacked, impairing stations look similar: with 205, 203, and 207 million mean impairing population for four-, three-, and two-stage auctions respectively compared to 192 million for singe-stage auctions. With VHF repacking, we also again observe that multi-stage auctions performed much worse than single-stage auctions, with four-stage simulations costing 1.25 times more and experiencing 1.10 times more value loss than single-stage auctions.

6.2.2 Early Stopping

The results just presented led us to ask: is there anything the FCC could have done to determine a clearing target endogenously while softening the resulting cost and efficiency losses? We propose a simple answer, which we call *early stopping*: for each candidate clearing target, conducting the forward auction before the corresponding reverse auction, and stopping each reverse auction as soon as its cost exceeds forward auction revenue. To see why this would help, consider again our initial example in Figure 3, but imagine that the auctioneer knows going into the first reverse auction stage that the forward auction revenue is some number less than V_A. Once station B freezes at a price of V_A, the provisional cost of the first stage exceeds the first stage's forward auction revenue. Following our proposal, the first clock stage would immediately stop and so station C would not exit. Then, in the second stage, station B would exit next and the outcome would exactly match that of an auction in which the clearing target had initially been set to the final target. In general, continuing the reverse auction instead of aborting can only result in more stations exiting, reducing the possible choices for later stages.[12]

Early stopping does not always outperform the original design; it is possible to construct examples where the commitments made in earlier stages turn out to be better than the ones that would be made in later stages. A concrete example is provided in Appendix C. Nevertheless, our intuition was that early stopping would typically help in practice.

To test this, we ran two sets of experiments. The first set compared early stopping auctions against single-stage stations that "knew" the correct clearing target (similar to the experiments previously described for multi-stage auctions). The second set of experiments compared the amount of spectrum cleared by early stopping auctions to auctions using the original design.

Both of these experiments required forward auction revenues as input: this is the part of our simulations for which data was thinnest. We adopted a convenient model described in Appendix E. We used the same forward auction revenues across all sampled station value profiles for a given value model.[13] Since we had explicit forward auction values and known clearing costs, the auctions could end at any stage, meaning that the amount of spectrum being sold could differ even across paired simulations. We assume that clearing more spectrum is preferable to clearing less, noting the FCC's stated goals in running the auction.

We implemented the early stopping rule as follows. Whenever a station became a provisional winner, we checked to see if the cost exceeded the forward auction revenue for that stage. If it did, then we designated all remaining bidding stations as provisional winners at their current prices, and the reverse auction continued on to the next stage as normal.[14]

We began our experiments by running early stopping auctions with an initial clearing target of 126 MHz (corresponding to the first stage of the real auction) and following the 600 MHz band plan (see Figure 7) from then on. Once these experiments completed, we determined the final stage of each early stopping simulation and ran a corresponding single-stage auction in order to assess the "penalty" due to early stopping vs. perfect forecasting of the clearing target. The results are shown in Figure 10. For computational reasons, we only performed experiments for UHF-only simulations (the experiments took more than 15 CPU years overall). To avoid making comparisons between different qualities of spectrum, as described above, we disabled our impairment mechanism and let prices start as high as required for full UHF participation for these experiments. We observed that in both value models, early stopping performed well relative to the single-stage auctions, with increases to average cost and value loss of no more than about 1%. These results are particularly encouraging when compared to the multi-stage experiments described earlier, in which we observed significant gaps between multi-stage clearing and perfect forecasting.

We next ran simulations that compared early stopping to the original incentive auction to determine which cleared more spectrum. Under both of our value models, on average more mobile licenses were created when using early stopping over the original algorithm—0.50 and 0.22 extra licenses under the BD and MCS models, respectively. Early

[12] A suggestion along similar lines was given by Ausubel et al. (2017). Rather than swapping the order of the reverse and forward auctions, they proposed forecasting reasonable bounds on forward auction revenue (i.e., asserting that it would be unlikely for telecoms to pay more than $X for a given amount of spectrum) for each stage. The reverse auction would then terminate when the provisional cost reached $X, the following forward auction would be skipped, and the next stage of the reverse auction would be begin. They argue that first, this would have reduced the auction duration, and second, that bidders in the forward auction had a sense in early stages that the prices were too high for the auction to terminate, so they did not bid sincerely.

[13] We found that under the MCS value model, early stopping auctions almost always terminated in the first stage; since this led to uninformative experiments, we halved the forward auction revenues described in the appendix when using the MCS value model.

[14] We made one other change. We noted that the reverse auction has a step at the end of each round where it removes unconstrained stations: those whose clocks will provably wind down to zero in a given stage and therefore might as well exit as soon as this is detected. In an early stopping auction, this no longer makes sense, as such stations may not be unconstrained in the next stage. Therefore, we did not perform these checks in our early stopping auctions.

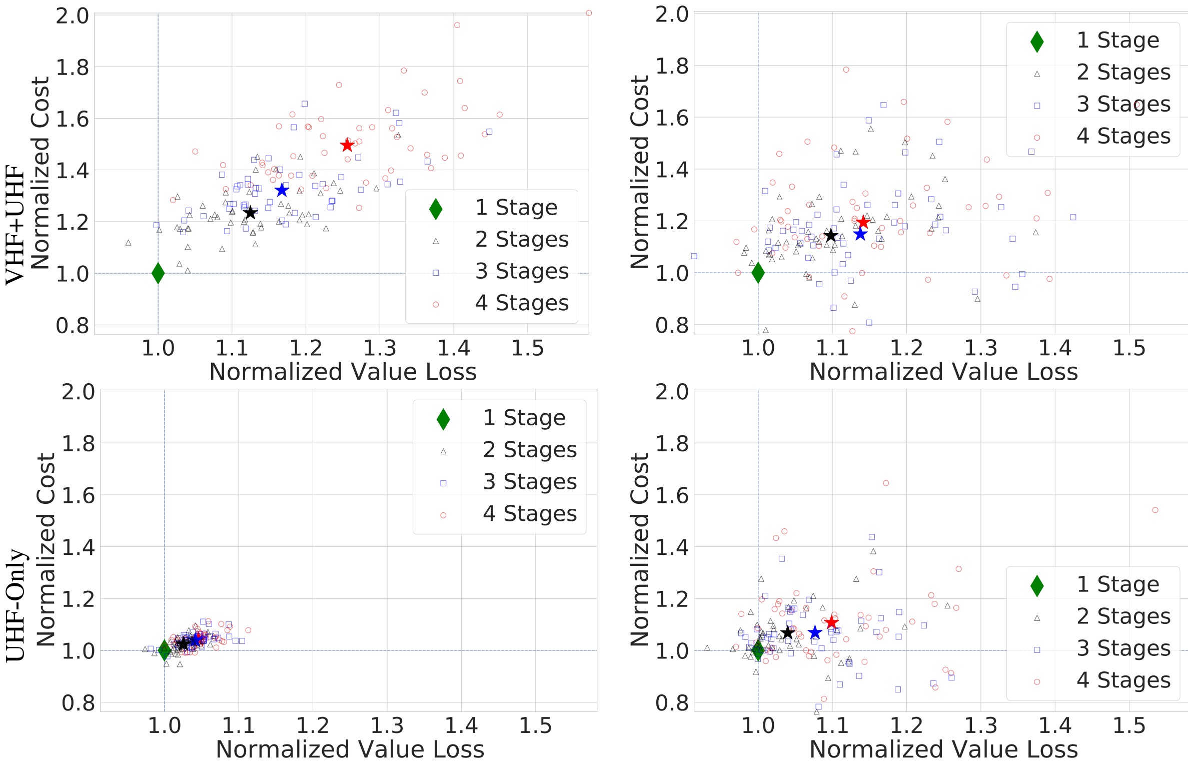

(a) MCS

(b) BD

Figure 8: Comparing auctions running through 1-4 stages, ultimately ending on the same clearing target, with no impairing stations.

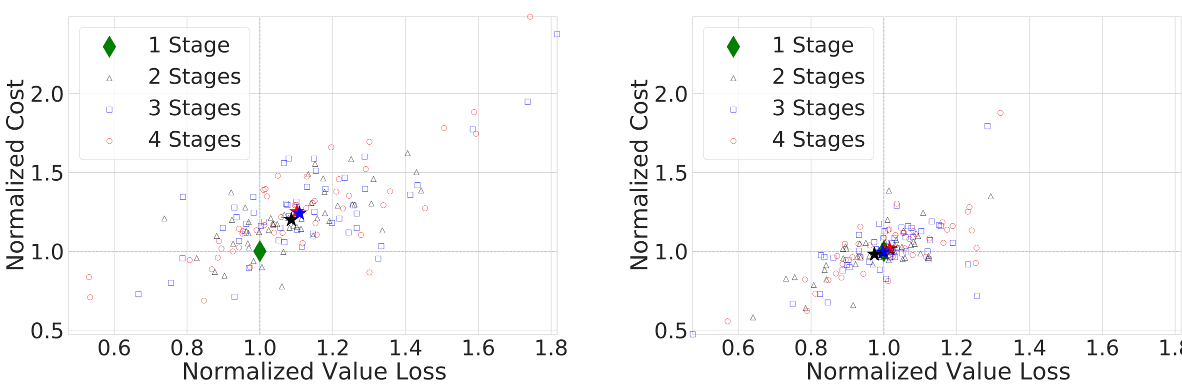

(a) VHF+UHF

(b) UHF-Only

Figure 9: Comparing auctions running through 1-4 stages, ultimately ending on the same clearing target, factoring in impairing stations. The BD value model is used for all simulations.

stopping also led to shorter auctions: on average early stopping simulations required only 30% and 76% of the number of rounds required by the original incentive auction under the BD and MCS models, respectively.

We conclude by reflecting on some concrete numbers from the real incentive auction. At the end of the very first round of the first stage, payments to frozen stations already exceeded \$50 billion; when the stage finished, these payments were \$86 billion. In the subsequent forward auction, revenues were only \$23 billion. Early stopping would have terminated the reverse auction during the very first round of bidding. The first reverse auction stage took one full month to resolve. This month, and possibly more time in future stages, would have been saved if early stopping had been implemented.

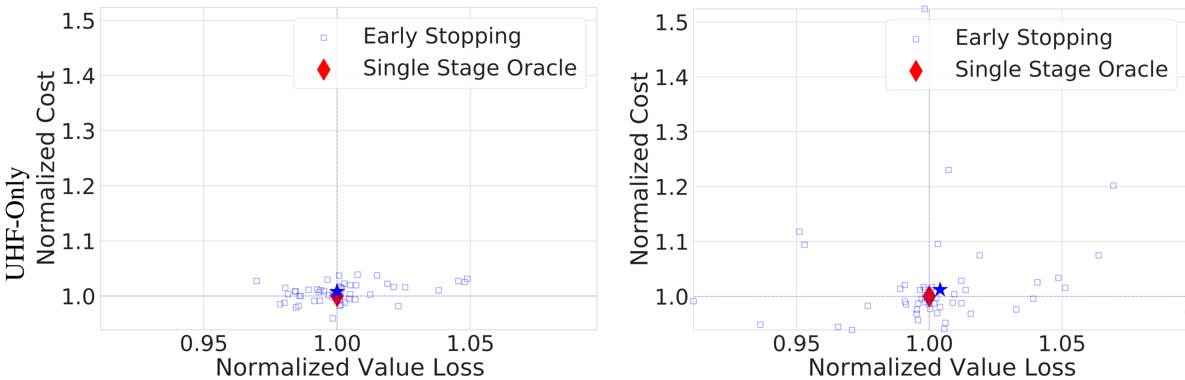

(a) MCS

(b) BD

Figure 10: Comparing auctions using the early stopping algorithm against single-stage auctions ending on the same clearing target as their corresponding early stopping auction.

6.3 Scoring Rules

Stations' starting prices in the incentive auction were not all the same. Instead, they were set proportionally to an assigned *score*, determined by a *scoring rule*. For active stations, the initial relative prices for going off-air remained constant during the auction, so those prices impacted the order in which stations exited the auction and continued to broadcast.

Theoretically, scoring performs two distinct functions. First, because every descending clock auction is equivalent to a greedy algorithm for packing stations into the broadcast spectrum, it may be possible to pack a larger and more valuable set of stations if the algorithm prioritizes stations that create less interference with other stations. In the FCC's design, this idea was implemented by offering higher prices to stations with more "interference links"[15] so that those stations would be less likely to exit. The second function of scoring is to reduce the cost of the auction by offering lower prices to stations that would be likelier to accept them. In the FCC's design, this was implemented by reducing the clock prices offered to stations serving smaller populations. This element of the design was informed by the formal theory of how to set prices to minimize expected total costs developed by Myerson (1981).

Putting these two elements together, opening prices in the incentive auction were set in proportion to the square root of the product of a station's population and its interference links. The population component of scoring was one of the auction's most controversial design elements. Indeed, it was vigorously opposed by a coalition of owners of lower powered stations serving smaller populations (the Equal Opportunity for Broadcasters Coalition), whose starting prices in the clock auction were reduced by this scoring rule.

In our experiments, we compared the following four scoring rules:[16] (1) "Incentive Auction" the scoring rule used in the actual auction (2) "Interference", the square root of a station's interference links (3), "Population", the square root of a station's population, and (4) "Uniform", a rule where each station has the same score. As in the actual auction, we normalized all scores so that the highest scoring UHF station had a score of 1 million.

We note a difficulty in these experiments: scoring rules impact prices, so our impairment mechanism will not necessarily select the same set of impairing stations for each scoring rule on the same value profile. To prevent "apples-to-oranges" comparisons across auctions clearing spectrum of varying quality, we disabled our impairment mechanism, instead starting the base clock price as high as was required (in multiples of 5% increases of the FCC's base clock price, to achieve the same reduction schedule) to ensure full UHF-station participation.[17] Another reason to not run

[15]More precisely: "an index of the number and significance of co- and adjacent channel interference constraints that station would impose on repacking."

[16]We used population and interference values from the FCC's opening prices document (FCC 2015d), noting that the links heuristic is calculated on the full interference graph. We continue to use these values in our UHF-only auctions, even though the underlying interference graph changes (the number of links goes down). We observe that the interference graph somewhat similarly changed between each stage of the incentive auction (the values correspond to the full graph, or the final stage of the band plan) but nevertheless the same values for the links heuristic were used throughout the auction.

[17]When repacking the VHF band, we continued to lock the VHF band until the base clock price reached the real FCC starting base clock price, because the heuristics that set VHF prices are fragile and we were concerned that altering them might cause unintended effects on the VHF allocation.

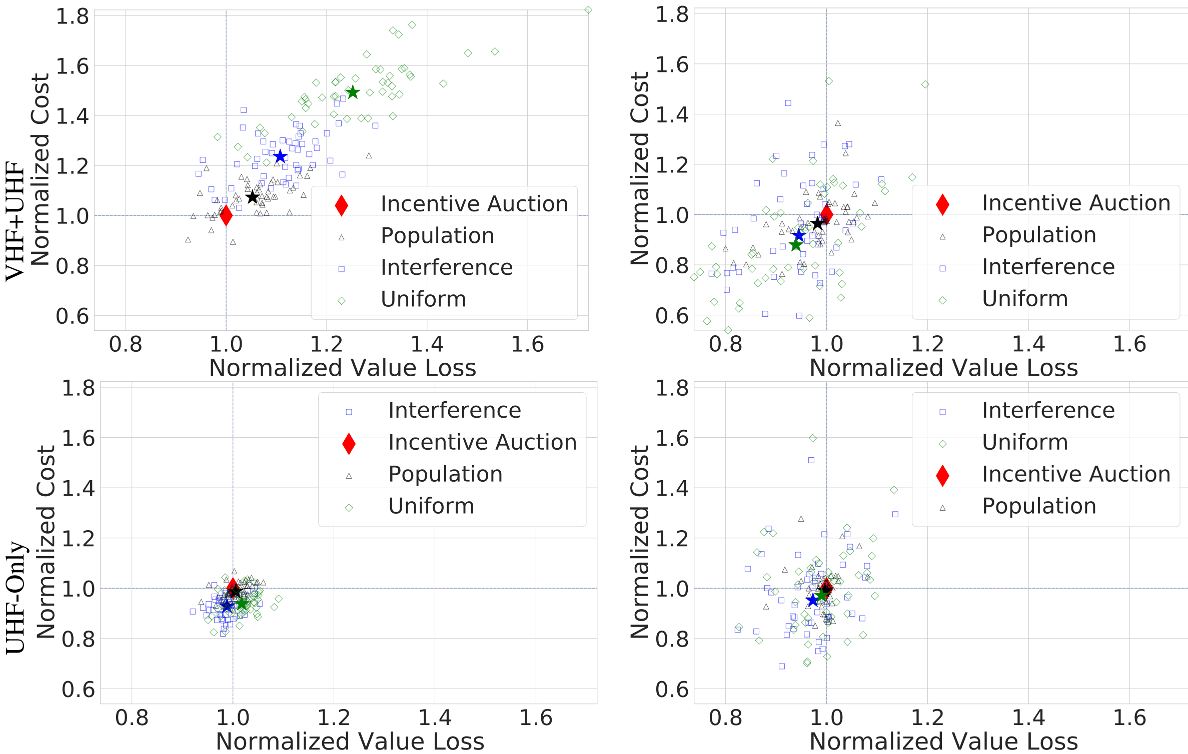

(a) MCS

(b) BD

Figure 11: Comparing auctions using four different scoring rules.

our impairment mechanism in this case is that the starting base clock price chosen by the FCC would surely change under a shift in the distribution of station scores; by not running our impairment procedure our results are not sensitive to this parameter. Running with no impairment mechanism corresponds to considering a world in which the FCC is unwilling to accept any degradation to the cleared spectrum.

Results for our simulations using the various scoring rules are shown in Figure 11. Under the MCS value model when the VHF band was repacked, we observed that simulations using only interference scoring cost 1.24 times as much and experienced 1.11 times as much value loss as simulations that combined population and interference scoring (the FCC's scoring rule). Uniformly scoring stations was not effective in this setting, with simulations averaging nearly 50% higher costs and 25% higher value loss. With the same value model, when only the UHF band was repacked, the scoring rule appeared to matter much less. Here, the interference scoring rule outperformed other scoring rules on average, with mean costs and value losses of 0.93 and 0.99 times respectively compared to corresponding simulations using the FCC's scoring rule. Under the BD value model, we observed much higher variance across simulations. If we nevertheless consider average performance, the FCC's scoring rule was outperformed in both metrics by every other scoring rule considered regardless of whether the VHF band was repacked. When the VHF band was repacked, the lowest value losses and costs were achieved, surprisingly, by the uniform scoring rule (94% mean value loss and 88% cost relative to the FCC's scoring rule). When only the UHF band was repacked, the lowest average value losses and costs were achieved by intereference scoring (97% and 95% of the FCC scoring rule). On the whole, outside of one setting (MCS values and UHF + VHF repacking) we did not find strong evidence for population-based scoring. These experiments required 3.5 years of CPU time to run.

All of the scoring rules that we considered are static: they assign each station a single, fixed score. The theory of descending clock auctions also considers scores that can be set to respond to the history of the auction in a dynamic fashion (provided that they never increase a station's price). We leave the investigation of dynamic scoring rules for this setting as future work, but note that our simulations' failure to crisply recommend any of the static scoring rules we considered may advocate for such an investigation.

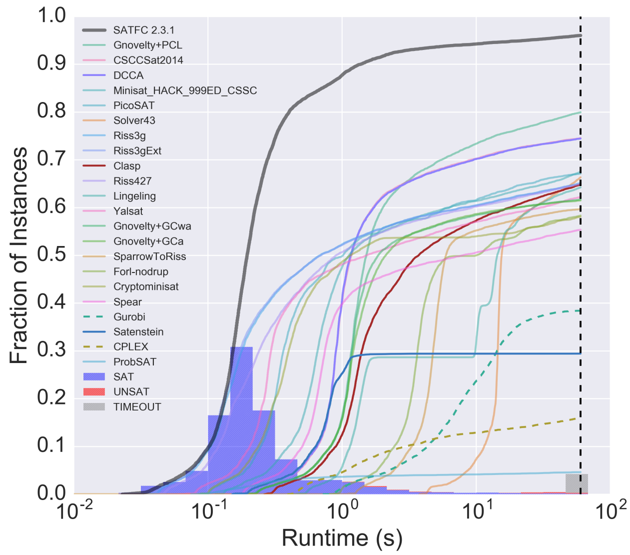

Figure 12: Empirical Cumulative Density Function of runtimes for default configurations of MIP and SAT solvers and for SATFC 2.3.1 on a benchmark set of 10 000 non-trivial problems. The curves show fraction of instances solved (y axis) within different amounts of time (x axis; note the log scale). The legend is ordered by percentage of problems solved before the cutoff. The histogram indicates density of SAT and UNSAT instances binned by their (fastest) runtimes; unsatisfiable instances constituted fewer than 1% of solved instances.

6.4 Feasibility Checking

The feasibility checker determines if a given station can be repacked alongside the stations that have already exited the auction. Feasibility checking was a large concern in the incentive auction because the station repacking problem is hard both theoretically (it is NP-complete) and in practice. When the feasibility checker cannot find a way to repack a station (either by proving that no repacking exists or by running out of time), a station freezes—it stops bidding and its compensation stops falling. The feasibility checker's quality therefore has a direct effect on both the cost and efficiency of the auction: if the feasibility checker fails to find an assignment that repacks a station when such an assignment exists, this station will freeze at an unnecessarily high price. If the station would otherwise never freeze at all, the value of the allocation is changed.

Do auction outcomes improve according to our efficiency and cost metrics as the feasibility checker improves? Our intuition is that if the auction algorithm is a good one, then better feasibility checking should lead to better results. However, this intuition is not always correct; see Appendix D for a counterexample.

6.4.1 Effect of the Feasibility Checker on Auction Outcomes.

SATFC 2.3.1, the feasibility checker that was used in the incentive auction, was designed over several years (Newman et al. 2017, Fréchette et al. 2016).[18] The solver combines complete and local-search SAT-encoded feasibility checking with a wide range of domain-specific techniques, such as constraint graph decomposition. The authors used automatic algorithm configuration techniques to construct a portfolio of eight complementary algorithms from these various components. We now investigate whether the effort invested in making SATFC 2.3.1 was helpful, or whether a more off-the-shelf solution would have sufficed. To do so, we ran simulations in which we exchanged SATFC 2.3.1 with various alternative solvers.

We ran 22 solvers from ACLib (Hutter et al. 2014) (a library of solvers that support algorithm configuration) on a benchmark set of sampled station repacking problems from Newman et al.'s (2017) reverse auction simulations.[19] The

[18]While the prototype of SATFC was developed quickly, various unrelated legal and logistical delays to the auction's launch provided SATFC's authors with time to refine it.

[19]This benchmark can be downloaded at `https://www.cs.ubc.ca/labs/beta/www-projects/SATFC/cacm_cnfs.tar.gz`

results are shown in Figure 12. We selected certain solvers from this figure and ran simulations, using each as the feasibility checker.[20] Specifically, we decided to compare the following solvers.

- SATFC 2.3.1: the feasibility checker used in the incentive auction.
- Greedy: a solver that simply checks whether a previous assignment can be augmented directly, without changing the assignments of any other stations. This algorithm is the simplest reasonable feasibility checker and thus serves as a baseline.
- PicoSAT: To our knowledge, alongside MIP approaches, PicoSAT is the only other solver that has been used in publications on the incentive auction, probably because it was shown to be the best among a set of alternatives in an early talk at the FCC on the subject (Leyton-Brown 2013).
- Gurobi, CPLEX: MIP solvers initially considered by the FCC.
- Gnovelty+PCL: the best performing of the 22 ACLib solvers on the benchmark data described above.

Our simulator always attempted to solve each feasibility check using the greedy algorithm before calling another solver, so every other feasibility checker can be understood as a sequential portfolio of the greedy solver and itself. We used this heuristic for good reason: the vast majority of feasibility checking problems encountered in a typical simulation can be solved greedily. Lastly, we note that for consistency we always used SATFC 2.3.1 as the feasibility checker in all experiments during the impairment phase (until the base clock hit the FCC's starting base clock price).

Our experiments took just over two CPU years. Results are shown in Figure 13. We observed that in general, stronger feasibility checkers led to better outcomes according to both of our metrics: the relative rankings of the solvers in the benchmark study translate exactly into the relative rankings across both of our metrics, regardless of whether the VHF band was repacked and regardless of our choice of value model. In particular, we observed that SATFC 2.3.1 dominated each other solver on both metrics, not only on average but across each individual simulation and across both value models. When compared to the best alternative off-the-shelf solver, reverse auctions based on *gnovelty+pcl* cost between 1.22 and 1.45 times more on average (depending on bands being repacked and the value model used) and lost between 1.22 and 1.45 times as much broadcaster value as those based on SATFC 2.3.1.[21]

6.4.2 Alternative Bid Processing Algorithms.

Given our findings about the impact of a strong feasibility checker, the reader might observe that the amount of time allocated to solving individual repacking problems could be important too. Could allowing more time for solving each problem have led to fewer stations freezing unnecessarily and overall improved economic outcomes? Figure 12 makes a compelling case that problems are overwhelmingly infeasible if they are not solved quickly, yet it is difficult to say without experimentation what the impact of solving the remaining truly feasible problems might be. Large increases in the cutoff time given to each problem would not have been practical. The first round of the reverse auction had 1030 bidding stations. A one hour cutoff, say, would have required more than 42 days in the worst case of sequentially checking stations. This would have slowed the pace of the auction to a halt!

We now explore two variants to the bid processing algorithm that increase the computation time available to processing stations before declaring them frozen while also respecting the time constraints of the auction. The first optimistically does not freeze a station with an indeterminate check; the second leverages parallel computing to utilize all of the computation time in a given round.

Revisiting Indeterminate Feasibility Checks. The incentive auction conservatively froze a station whenever the feasibility checker could not prove that the station could be repacked. We note that indeterminate cases did not have to be treated this way: what is important (for incentive purposes) is only that a station's winning price does not decrease if it cannot exit. While calling these stations infeasible saves computation time as there is no need to recheck them in later rounds, it is possible that as other stations exit, a given station's repacking problem becomes more constrained and perhaps easier to solve. An alternative bid processing algorithm, upon encountering an indeterminate feasibility check for a station, simply does not process that station's bid. Such a station is not frozen and will be asked to bid again in the next round. The prices offered to that station decreases as normal, but its winning payment $\mathcal{P}(s)$ will not decrease unless the solver finds a feasible repacking for the station. The auction designers rejected this alternative on the grounds that the actual rules would be simpler for bidders, but until now the alternative has never been quantitatively investigated.

[20] Similar experiments were performed in (Newman et al. 2017); the experiments reported here go beyond those by the inclusion of new feasibility checkers, UHF-only experiments, and results on a new value model.

[21] We emphasize that these experiments considered only each solver's default parameter setting, and do not investigate the performance each solver could have achieved if it had been tuned or otherwise customized.

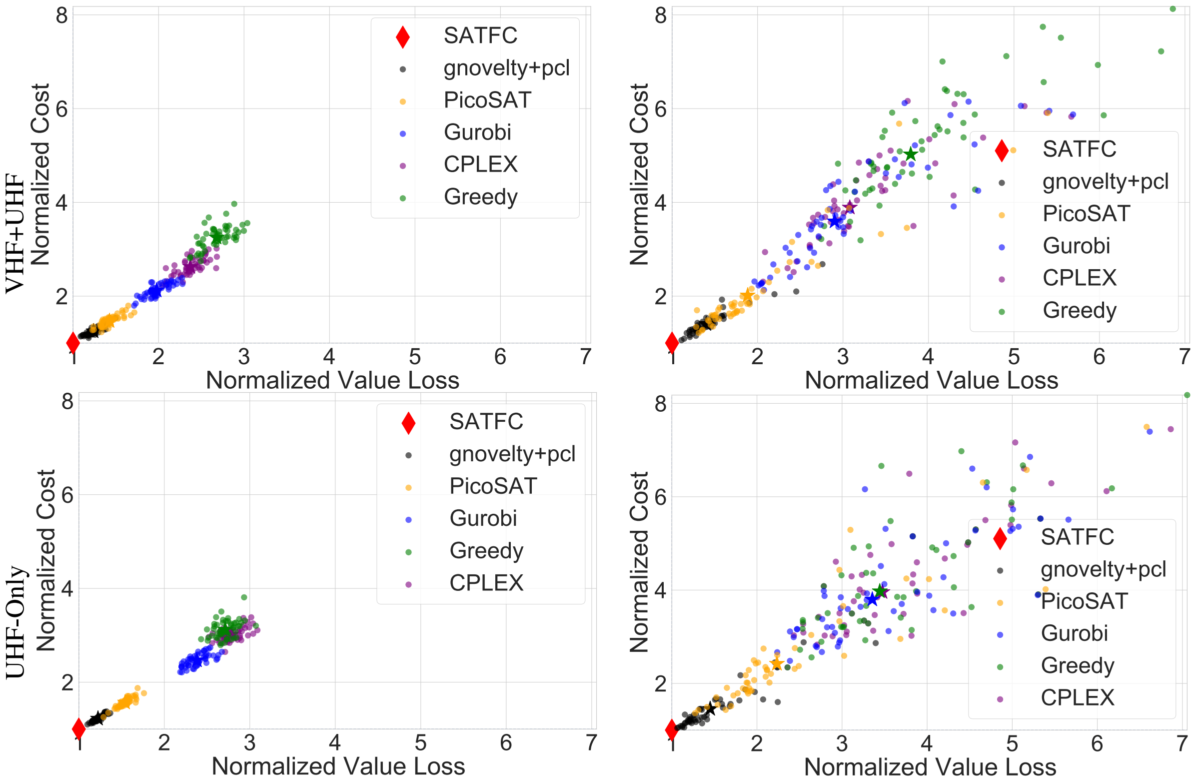

(a) MCS (b) BD

Figure 13: Comparing auctions using different feasibility checkers.

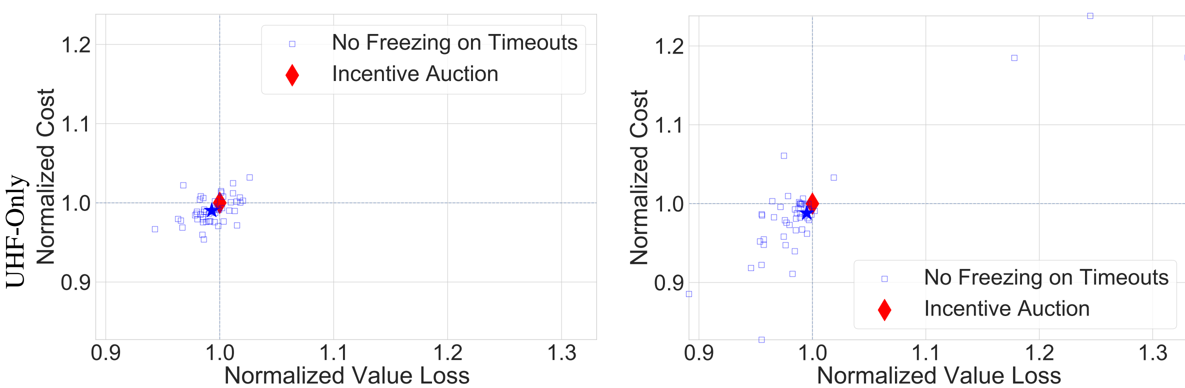

(a) MCS (b) BD

Figure 14: Comparing the standard bid processing algorithm with one that does not freeze stations with indeterminate feasibility checks.

Figure 14 compares the results of the standard bid processing algorithm against one that does not freeze stations for which the feasibility checker returns indeterminate results. These experiments took 3 CPU years to run. We observed on average, in UHF-only simulations, small (about 0.5–1%) improvements to both metrics when revisiting indeterminate results.

The First to Finish Algorithm. To ensure a predictable pace for the auction, bid processing for each round was required to complete within a fixed time window. Since stations were checked sequentially, with a fixed cutoff time allotted to each check, and since empirically most checks finished very quickly (see Figure 12), many rounds terminated earlier than they were required to. We wondered if a bid processing algorithm that wasted less computation time might have been able to solve more problems in each round.

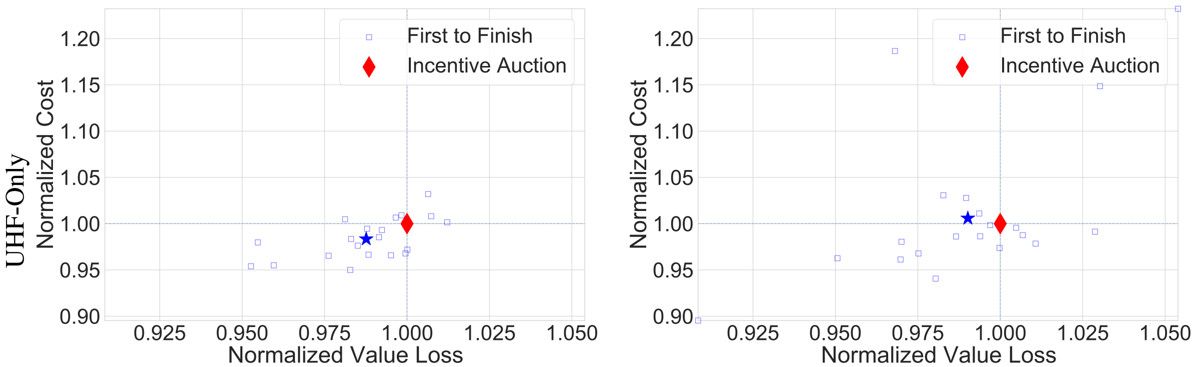

(a) MCS (b) BD

Figure 15: Comparing the first to finish algorithm against the standard bid processing algorithm for single-stage 126 MHz auctions.

Algorithm 2 First-to-Finish Bid Processing

```
Unprocessed ← S_bidding
while Time remains in the round and |Unprocessed| > do
    Launch a parallel feasibility check for each s ∈ Unprocessed, with a cutoff equal to time remaining in the round
    repeat
        Wait for a check to complete
        if s is not frozen then
            Process s's bid and remove s from Unprocessed
    until A station exits or moves bands
    Interrupt all ongoing checks
```

Based on this intuition, we propose an alternate bid processing algorithm, to which we refer as "first to finish".[22] The key idea is to run all of the feasibility checks in parallel,[23] with a cutoff equal to the entire amount of time remaining in the round.

As soon as any run completes, we process the corresponding station's bid, determining whether the station will move bands or exit the auction. If the station remains in its current band, we leave all other feasibility checks running unchanged; otherwise, we stop all of them, update the corresponding feasibility checking problems to include changes implied by the just-processed station, and start all parallel runs again with a new cutoff corresponding to all remaining time. We provide pseudocode in Algorithm 2.

The first-to-finish algorithm consumes dramatically more compute power than the standard bid processing algorithm: in the first round of the incentive auction, with 1030 stations to check, running this algorithm would have required on the order of 10 000 CPUs (noting that each run of SATFC 2.3.1 requires 8 parallel threads). However, at the scale of the incentive auction and given modern cloud computing resources, such computational requirements are not prohibitive.

Our experiments compare the first-to-finish algorithm[24] with a one-hour-per-round cutoff against the traditional bid processing algorithm with the usual one-minute-per-problem cutoff. Due to the computational requirements of these experiments, our results are based on only 20 samples; our experiments took more than 6 CPU years to run. The results are shown in Figure 15. We observed a fair amount of noise in both metrics due to varying the bid processing order and a small average effect: about a 1% improvement in both metrics under the MCS value model, and a 1% improvement to value loss and a 0.5% increase in cost under the BD value model. While these experiments were not

[22] In fact, we proposed this idea as the last details of the incentive auction were being finalized. Although the design team received the idea positively, it was not pursued because altering the bid processing algorithm would have contradicted the published auction rules, which was not possible by the time the idea emerged.

[23] The algorithm can be straightforwardly modified to make use of sequential runs with exponentially growing captimes in settings where it is impractical to conduct parallel feasibility checks for each station. In practice, a useful middle ground is to filter out easy problems by sequentially running every feasibility check with a short cutoff like 1 second, then to proceed with running the problems that remain in parallel.

[24] Since we lacked access to the required number of parallel computers, we ran a sequential simulation of the first to finish algorithm. The sequential version runs all of the unprocessed checks for a small cutoff and then processes them in order of completion time until encountering a movement or exit bid, at which point checks are restarted. If at the end of the cutoff, only indeterminate problems remain, the cutoff is set to the minimum of double the current cutoff or the remaining time left in the round.

exactly equivalent to increasing runtime cutoffs under the standard bid processing algorithm,[25] these results lead us to believe that increasing the cutoff time given to the feasibility checker would not have made a significant difference to the auction outcome. We do expect that the first-to-finish algorithm would have yielded larger gains if paired with a weaker feasibility checker or a larger incentive auction that gave rise to harder feasibility checking problems.

7 Conclusions

We used simulations to investigate previously unanswerable questions about the cost and efficiency of certain alternative designs for the incentive auction. To validate the robustness of our results, our simulations used two quite different value models: one from the empirical economics literature and another that we constructed to rationalize public bid data. At the scale of the incentive auction, even small percentage improvements in cost and welfare can translate into billions of dollars of savings.

Our main findings were that:

- repacking VHF led to significantly lower costs and more efficient outcomes;
- the multiple stage clearing rule substantially both increased costs and reduced the efficiency of the auction;
- a simple amendment to the clearing algorithm could both speed up the auction and nearly completely eliminates the multi-round inefficiency;
- the performance of pops scoring relative to other scoring rules varied widely based on the value model and whether the VHF band was repacked;
- the specialized feasibility checker developed for the auction significantly improved both cost and efficiency;
- alternative bid processing algorithms, while helpful, would not have made a significant difference to auction outcomes.

We hope these specific insights can help to inform future auction designs. More broadly, we believe our analysis demonstrates that large-scale statistical analysis of the simulated behavior of candidate market designs in highly complex settings—requiring substantial, but not unrealistic computational resources[26]—is a practical tool for understanding and evaluating alternative market designs.

A Robustness Experiments

When we described the BD value model (Section 4.1.2), we explained having made several design decisions. After performing our analysis, we wanted to understand how much our results depended on the specifics of those decisions. Here we describe the results of rerunning several experiments under four alternate parameterizations of our value model. These changes concern how we derive VHF valuations from UHF valuations. Recall that we model a UHF station's value for switching to the HVHF band as $\frac{2}{3} \cdot v_{s,\mathrm{UHF}} \cdot \mathcal{N}(1, 0.05)$ and similarly for the LVHF band with $\frac{1}{3}$ instead of $\frac{2}{3}$. These fractions were chosen for simplicity, drawing on some degree of domain knowledge about the values of VHF bidders. What if we had used other fractions instead? The FCC set opening off-air prices for HVHF and LVHF stations at $\frac{3}{5}$ and $\frac{1}{4}$ respectively of corresponding UHF stations. We reran experiments using $\frac{3}{5}$ and $\frac{1}{4}$ as the fractions in our value model. We refer to this parameterization as "Lower Values", as stations have lower values for the VHF bands when compared to our standard value model. For symmetry, we also reran experiments incrementing the fractions by corresponding amounts, leading to $\frac{11}{15}$ and $\frac{5}{12}$. We refer to this parameterization as "Higher Values". We also were interested in whether the Gaussian noise was consequential, so we created another parameterization where no noise was applied which we refer to as "No Gaussian Noise".

For each of our three value model parameterizations, we reran one experiment related to each question we posed in the introduction. Specifically, we compared: auctions that repacked the VHF band with those that did not ("Repacking VHF"); auctions that ran for four stages against those that simply ran for a single stage beginning at the fourth clearing target ("Clearing Procedure"); auctions that used the FCC's scoring rule against those that only used the population

[25] Prioritizing processing bids from stations whose checks are quick to solve may bias the ordering.

[26] Altogether, our final experiments took more than 60 CPU years to run. Writing this paper consumed perhaps 200–300 CPU years of compute in total, since we often found ourselves needing to rerun experiments as we uncovered bugs, changed parameters, or refined our experimental questions. We do not provide these CPU time numbers with the goal of impressing the reader with how much compute power we used. Instead, we hope (a) to give other researchers a sense of the scale of experiments that we consider necessary for answering questions like the ones we tackled here, and (b) to reassure the reader that we ran more or less the largest experiments that were practically feasible.

component ("Scoring Rules"); and lastly, auctions that used SATFC 2.3.1 as their feasibility checker against those that used the best off-the-shelf feasibility checker ("Feasibility Checker"). We ran 50 paired simulations for each experiment, just as we did for the corresponding main results. These experiments took roughly 20 CPU years to run.

The results are summarized in Table 1. We observed slightly different results among the three parameterizations: for example, the cost savings estimate from repacking VHF bands varied from 14–26% depending on how substitutable stations felt the VHF bands were for the UHF band. However, no change to the value model substantially altered the conclusions we drew from any experiment, giving us confidence that our results are relatively robust to the particular choices we made.

Model Change	Value Loss Repacking VHF	Clearing Procedure	Scoring Rules	Feasibility Checker
None	0.92 (0.15)	0.95 (0.24)	0.94 (0.08)	1.44 (0.31)
Lower Values	0.96 (0.14)	0.95 (0.18)	0.97 (0.08)	1.51 (0.41)
Higher Values	0.87 (0.15)	0.91 (0.17)	0.93 (0.09)	1.53 (0.33)
No Gaussian Noise	0.9 (0.13)	0.94 (0.19)	0.95 (0.08)	1.48 (0.31)

Model Change	Cost Repacking VHF	Clearing Procedure	Scoring Rules	Feasibility Checker
None	0.78 (0.15)	0.85 (0.2)	0.96 (0.19)	1.44 (0.29)
Lower Values	0.86 (0.13)	0.89 (0.21)	0.98 (0.18)	1.48 (0.32)
Higher Values	0.74 (0.16)	0.81 (0.17)	0.93 (0.18)	1.55 (0.34)
No Gaussian Noise	0.79 (0.14)	0.85 (0.18)	0.95 (0.19)	1.49 (0.31)

Table 1: A summary of the results of changing the value model on four replicated experiments from the main paper. Each row corresponds to a different change to the value model, and each column corresponds to an experiment. The "None" row refers to the original experiment. Values in each cell represent the mean and standard deviation of value loss (upper table) and cost (lower table) of an altered auction design relative to the real auction design across all paired simulations.

B Multi-stage Auction Performance Counterexample

Example B.1 *Consider a setting with four identically scored stations A, B, C, D with $V_A > V_B > V_C > V_D$, $V_B < V_C + V_D$ and $V_A < 2V_B$. In the final stage, the feasible sets are $\{A, B\}, \{A, C, D\}$ and all subsets of these sets. As prices drop, A exits first, followed by B, which freezes C and D at prices of V_B. So the single-stage approach gives a value loss of $V_C + V_D$ and a cost of $2V_B$*

In the multi-stage setting, in the first stage assume the feasible sets are $\{B\}, \{A, C, D\}$ and all subsets of these sets. A will exit first, freezing B at a price of V_A. B and C then exit, concluding the stage. In the second stage, B never unfreezes. Therefore, the total cost of the multi-stage approach is V_A and the value loss is V_B. By the inequalities assumed above, this is a cheaper, more efficient outcome than the single-stage auction that predetermined the amount of spectrum to clear.

C Early Stopping Counterexample

Example C.1 *Consider four identically scored stations A, B, C, D with $V_A > V_B > V_C > V_D$, $V_B < V_C + V_D$ and $V_A < 2V_B$. Let the forward auction run first and have a purchasing price of V_A. Let the feasible sets in the first stage be $\{A, C\}, \{A, D\}, \{B\}$ and all subsets of these sets. In the second stage, the feasible sets add $\{A, B\}, \{A, C, D\}$ and all subsets. In both cases, the auction begins with A exiting and B freezing at price V_A. In an early stopping auction, this will trigger the end of stage one. In stage two, B unfreezes and exits. Then C and D freeze at price V_B. This leads to a value loss of $V_C + V_D$ and a cost of $2V_B$. In an auction without early stopping, C would exit, freezing D at price V_C. This would trigger the stage to end. In the next stage, B would remain frozen and D would unfreeze and exit, leading to a value loss of V_B and a cost of V_A. Using the inequalities on the values above, the auction that does not use early stopping performs better in both metrics.*

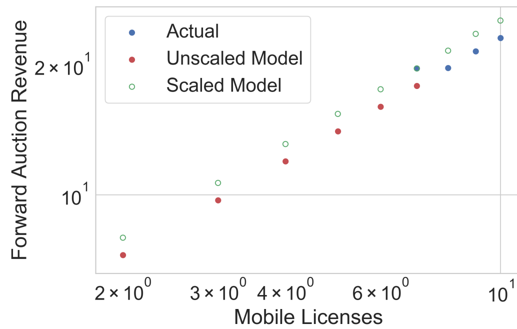

Figure 16: Forward auction revenues observed in the incentive auction ("Actual") and our model of forward auction revenue.

D Feasibility Checker Counterexample

For a fixed cutoff, a feasibility checker can be thought of as a mapping from a set of stations to $\{Feasible, Infeasible, Unknown\}$. We can define an ordering over feasibility checkers such that a feasibility checker F_1 is strictly better than a second F_2 if and only if for all possible sets of stations $s \in 2^S$, $F_2(s) = Feasible \implies F_1(s) = Feasible$ and $\exists s$ such that $F_1(s) = Feasible$ and $F_2(s) = Unknown$. Importantly note that for the purposes of this definition we don't care about the feasibility checker's ability to prove infeasibility: while this is important for saving time and for understanding whether there is room to improve existing feasibility checkers, it does not ultimately impact the result of the auction since infeasibility and indeterminate solutions are treated identically.

Example D.1 *Imagine a UHF-only setting involving four bidding stations A, B, C, and D. Let $V_A = V_B = V_C = V_D = V$ and let all stations have the same score. The constraints are such that the repackable sets are either $\{B, C, D\}$ or $\{A, D\}$ (and all subsets). There are two feasibility checkers: F_1 can find all feasible repackings, but $F_2(\{A, D\}) =$ Unknown. Consider the first round in which each station is being offered a price p just below V and assume that the bid processing order is D, A, B, C. D exits the auction. Under F_1, A is allowed to exit the auction. This freezes B and C. The value loss for these two stations will be $2V$ and the payment will be just under $2V$. F_2, however, cannot pack A, and so A freezes. B and C then exit the auction. The value loss in this scenario is V and the payment is just under V, so this outcome is strictly better than the previous even though F_1 is strictly better than F_2.*

E Modeling Forward Auction Revenue

When using the early stopping algorithm, the reverse auction takes as input the output of the previous forward auction (the amount that mobile carriers will pay for the spectrum). Therefore, in order to simulate early stopping auctions, we need to model forward auction revenues.

While we have access to the real forward auction revenues, the incentive auction only went through four stages, and simulations could potentially go to stages beyond the fourth. To address this issue, we performed a log-log fit on the number of mobile licenses and forward auction proceeds in the first three stages (i.e., with 3 data points we fit $ln(\#\text{licenses}) = a \cdot ln(cost) + b$ for some constants a, b). We ignored the revenue in the fourth stage when performing this fit because the price per license rose significantly relative to the other stages, and we suspect that the price increase was likely due to an understanding among bidders that the auction would terminate in this stage. After the fit was established, we scaled the entire model by a constant so that the model's prediction for the fourth stage matched the observed real revenue. The results appear in Figure 16.

A final problem is that when using the MCS model, we observed that reverse auction simulations rarely produced high enough procurement costs to trigger early stopping in the first stage, leading to uninformative behavior in which the auction simply ends after the first stage. To get around this, we scaled our forward auction model downwards by a factor of 2 when running simulations for the MCS.

References

Ausubel LM, Aperjis C, Baranov O (2017) Market design and the FCC Incentive Auction, Working Paper.

Bichler M, Fux V, Goeree JK (2019) Designing combinatorial exchanges for the reallocation of resource rights. *Proceedings of the National Academy of Sciences* 116(3):786–791.

Bichler M, Hao Z, Littmann R, Waldherr S (2020) Strategyproof auction mechanisms for network procurement. *OR Spectrum* ISSN 1436-6304, URL `http://dx.doi.org/10.1007/s00291-020-00597-7`.

Cramton P, López H, Malec D, Sujarittanonta P (2015) Design of the reverse auction in the FCC Incentive Auction.

Derakhshan M, Golrezaei N, Paes Leme R (2019) LP-based approximation for personalized reserve prices. *Proceedings of the 2019 ACM Conference on Economics and Computation*, 589–589.

Derakhshan M, Pennock DM, Slivkins A (2020) Beating greedy for approximating reserve prices in multi-unit VCG auctions. *arXiv preprint arXiv:2007.12653* .

Diamond S, Boyd S (2016) CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research* 17(83):1–5.

Domahidi A, Chu E, Boyd S (2013) ECOS: An SOCP solver for embedded systems. *European Control Conference (ECC)*, 3071–3076.

Doraszelski U, Seim K, Sinkinson M, Wang P (2017) Ownership concentration and strategic supply reduction. Working Paper, `http://www.nber.org/papers/w23034` 23034, National Bureau of Economic Research, URL `http://dx.doi.org/10.3386/w23034`.

FCC (2013) Office of engineering and technology releases and seeks comment on updated OET-69 software. *FCC Public Notice* DA 13-138.

FCC (2015a) Application procedures for broadcast Incentive Auction scheduled to begin on March 29, 2016. *FCC Public Notice* DA 15-1183.

FCC (2015b) Incentive Auction task force releases initial cleairing target optimization simulations. *FCC Public Notice* DA 15-606.

FCC (2015c) Repacking constraint files. `http://data.{FCC}.gov/download/incentive-auctions/Constraint_Files/`, accessed 2015-11-20.

FCC (2015d) Reverse auction opening prices. `http://wireless.{FCC}.gov/auctions/incentive-auctions/Reverse_Auction_Opening_Prices_111215.xlsx`, accessed 2015-11-15.

FCC (2019) Incentive Auction: Reverse auction - bids. `https://auctiondata.{FCC}.gov/public/projects/1000/reports/reverse-bids`, accessed 2019-11-28.

Fréchette A, Newman N, Leyton-Brown K (2016) Solving the station repacking problem. *AAAI Conference on Artificial Intelligence*, 702–709, URL `http://dl.acm.org/citation.cfm?id=3015812.3015917`.

Gkatzelis V, Markakis EK, Roughgarden T (2017) Deferred-acceptance auctions for multiple levels of service. *Proceedings of the 2017 ACM Conference on Economics and Computation* .

Golrezaei N, Lin M, Mirrokni V, Nazerzadeh H (2017) Boosted second price auctions: Revenue optimization for heterogeneous bidders. *Available at SSRN 3016465* .

Hutter F, López-Ibáñez M, Fawcett C, Lindauer M, Hoos HH, Leyton-Brown K, Stützle T (2014) AClib: A benchmark library for algorithm configuration. *Conference on Learning and Intelligent Optimization (LION)*, 36–40 (Springer).

Kearns M, Dworkin L (2014) A Computational Study of Feasible Repackings in the FCC Incentive Auctions. *arXiv:1406.4837 [cs]* URL `http://arxiv.org/abs/1406.4837`, arXiv: 1406.4837.

Leyton-Brown K (2013) The viability of exact feasibility checking, URL `https://www.cs.ubc.ca/~kevinlb/talks/2013-Feasibility-Testing.pdf`, talk at the Stanford Institute for Economic Policy Research Conference on the design of the U.S. Incentive Auction for reallocating spectrum between wireless telecommunications and television broadcasting.

Leyton-Brown K, Milgrom P, Segal I (2017) Economics and computer science of a radio spectrum reallocation. *Proceedings of the National Academy of Sciences* 114(28):7202–7209, ISSN 0027-8424, URL `http://dx.doi.org/10.1073/pnas.1701997114`.

Li S (2017) Obviously strategy-proof mechanisms. *American Economic Review* 107(11):3257–87.

Milgrom P, Segal I (2020) Clock auctions and radio spectrum reallocation. *Journal of Political Economy* 128(1):1–31, URL `http://dx.doi.org/10.1086/704074`.

Morgenstern JH, Roughgarden T (2015) On the pseudo-dimension of nearly optimal auctions. *Advances in Neural Information Processing Systems* 28:136–144.

Myerson RB (1981) Optimal auction design. *Mathematics of operations research* 6(1):58–73.

Newman N, Fréchette A, Leyton-Brown K (2017) Deep optimization for spectrum repacking. *Communications of the ACM* 61(1):97–104.

Nguyen TD, Sandholm T (2014) Optimizing prices in descending clock auctions. *Proceedings of the fifteenth ACM conference on Economics and computation*, 93–110 (ACM).

Nielsen (2022) DMA regions. `https://www.nielsen.com/us/en/contact-us/intl-campaigns/dma-maps/`, accessed 2022-03-06.

Paes Leme R, Pal M, Vassilvitskii S (2016) A field guide to personalized reserve prices. *Proceedings of the 25th international conference on world wide web*, 1093–1102.

Roughgarden T, Wang JR (2019) Minimizing regret with multiple reserves. *ACM Transactions on Economics and Computation (TEAC)* 7(3):1–18.

Wellman MP, Greenwald A, Stone P (2007) *Autonomous bidding agents: Strategies and lessons from the trading agent competition* (MIT Press).