ARTIFICIAL INTELLIGENCE AND MARKET DESIGN: LESSONS LEARNED FROM RADIO SPECTRUM REALLOCATION

A PREPRINT

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

Paul Milgrom
Department of Economics
Stanford University
Stanford, CA94305
milgrom@stanford.edu

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

Ilya Segal
Department of Economics
Stanford University
Stanford, CA94305
isegal@stanford.edu

1 Introduction

The modern AI revolution has been fueled by computational processes whose performance improves as they are allocated more data and more compute power. As computers have become faster and cheaper, and as data sets have grown larger, algorithmic approaches ranging from self-supervised neural networks to Monte Carlo Tree Search have exhibited skyrocketing performance. Unfortunately, most mainstream market design approaches exhibit neither of these desirable scaling properties: purely theoretical analyses are unaffected by Moore's Law; rationality assumptions often make observations of actual bidder behavior irrelevant beyond the (worthwhile but narrow) question of estimating valuations via econometric techniques. How might we go further, leveraging cheap compute power and abundant data to complement the work of auction analysts? One path forward is running extensive simulations that leverage a high-fidelity market simulator and realistic models of participant preferences. A (mostly) recent literature has applied this approach across a diverse variety of domains, each with their own literature, including matching mechanisms for refugee resettlement [Delacrétaz et al., 2019]; school assignment policies Allman et al. [2022]; dock allocation and rebalancing in bike sharing networks [Freund et al., 2018]; patrolling strategies for wildlife protection [Yang et al., 2014]; reallocating fishing licenses [Bichler et al., 2019]; and matching in organ exchanges [Santos et al., 2017].

This chapter[1]—based on a talk given at the National Bureau of Economic Research Conference on New Directions in Market Design in May, 2023—applies a qualitatively similar approach to the investigation of the Incentive Auction, a recent radio spectrum reallocation conducted by the Federal Communications Commission (FCC). There are many open questions in this setting: the auction's design was both novel and extremely complex [Leyton-Brown et al., 2017], and as a result it was not possible to thoroughly consider every potential design variation before the auction was run. Furthermore, the complexity of the setting presents a significant barrier to purely theoretical analyses. Our goal in this chapter is to show how a computationally intensive approach can be used to evaluate 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. Such insights can inform other important resource allocation problems that may leverage parts of the auction's design. More broadly, we hope this work will serve as an example for how AI can be employed to understand and evaluate alternative market designs in complex settings.

[1]This chapter substantially follows the paper "Incentive Auction Design Alternatives: A Simulation Study" by the same authors, which is forthcoming in *Management Science: Revenue Management and Market Analytics*. Please refer to that paper for details regarding the setup of the experiments described here and for additional experimental results. The chapter also draws in places from "Deep Optimization for Spectrum Repacking", N. Newman, A. Fréchette, K. Leyton-Brown, *Communications of the ACM*, volume 61, number 1, pp. 97–104, January 2018.

1.1 Evaluating Complex Auction Designs

To date, the design and analysis of auctions has relied primarily on theoretical tools. Such analysis becomes increasingly difficult as auctions become more complex[2], e.g., as the number of goods at auction increases; as the number of parameters needed to describe valuation functions grows; as heterogeneity across bidders increases; etc. Nevertheless, the incentive auction's design was strongly informed by theoretical analysis, which has shown that deferred acceptance (DA) auctions have many good properties. DA auctions exhibit "unconditional winner privacy" [Milgrom and Segal, 2020], meaning that winners are only required to reveal as much about their valuation as is necessary to prove that they are winners. When stations' bids are binary responses to a series of descending offers and when stations are all independently owned, DA auctions are obviously strategy-proof [Li, 2017] and weakly group strategy-proof. When the sets of stations that can be jointly repacked form a uniform matroid, DA auctions repack the efficient set of stations [Bikhchandani et al., 2011]; if furthermore bidder values are drawn independently from known distributions, scoring offers according to virtual values can implement the Myerson "optimal auction" [Milgrom and Segal, 2020]. In exchange for the universality of its findings, theoretical analysis necessarily relies on simplifications. Such simplifications can raise questions about the practical applicability of results.

A second approach for evaluating auction designs is data driven. There is a vast literature on laboratory and field experiments for auctions [Ferejohn et al., 1979, Rassenti et al., 1982, Kwasnica et al., 2005, Cason et al., 2011, Adomavicius et al., 2012]. To our knowledge no such experimental data exists evaluating the incentive auction; indeed, it would be very challenging to conduct realistic experiments, given the real auction's long duration and extremely large number of participants. Data can also be obtained from an auction's practical deployment. There is an extensive structural estimation literature in auctions and matching markets. Such analysis has lead to significant insights in data-rich domains such as school choice [Agarwal and Somaini, 2018, Calsamiglia et al., 2020], highway procurement auctions [Krasnokutskaya and Seim, 2011, Somaini, 2020], timber auctions [Athey et al., 2011], and ad auctions [Athey and Nekipelov, 2010]. Unfortunately, only a single auction has been held to date using the incentive auction design, leaving us with only a single sample of past bids. We nevertheless leveraged this single datapoint to inform a valuation model, described below.

This chapter focuses on a third approach, computational simulations. This approach can verify whether theoretical findings from simplified models carry over to more complex, real-world settings and does not rely on the availability of rich historical data. Simulations are particularly well suited to studying market settings with complex clearing mechanisms (because such complexity tends to preclude "cleaner" analytical results) and uncontroversial models of agent behavior (because this reduces the risk that the model upon which the simulation relies will produce unrealistic behavior). The FCC conducted its own (mostly unpublished) internal simulations, leveraging a variety of techniques from operations research [Kiddoo et al., 2019].

The broader research community conducted simulation studies of the reverse auction at various points in the design process and differing in the fidelity with which they modeled the auction mechanism; the degree to which they made simplifications for computational reasons; and the valuation models they employed. Before the auction mechanism was even finalized, Kearns and Dworkin [2014] used simulations to speculate on the amount of spectrum that might be reallocated. This work used a bidder model that consisted entirely of determining which bidders would participate, which they studied both using independent coin flips for every bidder and more complex models in which affiliated bidders made correlated decisions. Later in the design process, Cramton et al. [2015] used simulations to lobby for design changes such as changing the methodology for setting opening prices. They leveraged 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". After the auction concluded, Doraszelski et al. [2017] used simulations to estimate how profitable and how risky bidder collusion strategies might have been, restricting experiments to regional scale and furthermore employing approximations ("limited repacking") to speed up computation. They concluded that the auction could have mitigated the harm imposed by collusion by restricting the participation of affiliated bidders in the auction. This work leveraged a novel value model for bidders, which we discuss further in Section 3.1.1 because we used it in our own simulations. The simulator used in this chapter more accurately simulates the actual reverse auction design than any other simulator of which we are aware.[3] Finally, Ausubel et al. [2017] performed a post-mortem analysis of the incentive auction in a similar vein to this chapter, but with a primary focus on the forward auction; we discuss one of their proposed amendments to the clearing algorithm in Section 5.4.2.

[2]For example, Kelly and Steinberg [2000] note of the Progressive Adaptive User Selection Environment (PAUSE) auction format that their design is "probably too complex to admit much theoretical analysis".

[3]We make no comment about the fidelity of the FCC's own simulator, since details about it are not publicly available.

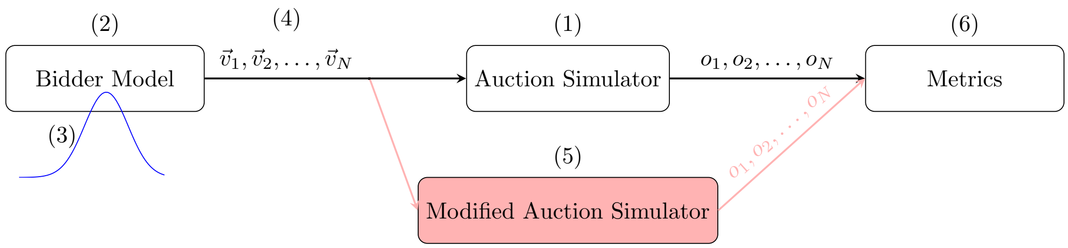

Figure 1: An overview of our simulation approach. N samples from a bidder model are fed into both a modified and unmodified simulator. The two outcome sets are converted into metrics and compared on a per-sample basis.

1.2 Our Simulation Methodology

We advocate and employed the following simulation methodology (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; and
6. Compare outcomes across paired samples using predetermined metrics.

For the study described in this chapter we instantiated these steps as follows. First, we custom-built the simulator used in this chapter. Other incentive auction simulators have been used in previous work (surveyed below) but ours is the most comprehensive of which we are aware, both in terms of its scale (national) and its coverage of the auction rules (e.g., multi-stage auctions including VHF bands). Our simulator's code is freely available online at `https://github.com/newmanne/SATFC`.

Next, perhaps the most critical step in our methodology is the specification of a bidder model, describing both valuations and behavior for each bidder. Since we focus on the reverse auction, each of our bidders corresponds to a television station. We considered two different valuation models: (1) the only fully specified model from the literature of which we are aware; and (2) a model we created based on publicly released bid data. We ran simulations with both models and contrast the results to investigate the robustness of our findings. Realistically modeling agent behavior is perhaps easiest in settings giving rise to *obviously dominant* strategies [Li, 2017]. A strategy is obviously dominant if any time an agent might deviate from the strategy, the best it can do under the deviation is no better than the worst it can do by following the strategy. Descending clock auctions make truthful bidding obviously dominant for bidders who own a single broadcast station (but not for those who own multiple stations).[4] Unfortunately, the real reverse auction lacked obviously dominant strategies because stations could do more than just accepting and rejecting an offer: they could also accept a lower price and downgrade to a VHF channel. We decided to study the problem by considering the assumptions both that VHF bidding could be abstracted away and that it could not. For the latter case, no equilibrium bidding strategy is known, let alone any model of how bidders would behave if they believed that some or all of their opponents were bidding out of equilibrium. Given the widespread belief that many bidders in this auction were unsophisticated, we modeled bidders as bidding truthfully and myopically (see Section 3.2).

Finally, we restricted most of our comparisons to simulations that cleared the same amount of spectrum, considering two key metrics: the sum of values of stations removed from the airwaves (value loss) and the aggregate amount paid to stations (cost). When assessing design elements that did affect the amount of spectrum cleared, we assumed that it was preferable to clear more spectrum, based on statements made by the FCC about the auction's intended goals [FCC, 2012].

[4] An agent can deviate from truthful bidding in this setting either by (1) accepting an offer below its value, or (2) by rejecting an offer above its value. Since prices only descend, these deviations either lead to selling at a loss or not selling at all, thus having a best-case utility of zero. Since truthful bidding has a worst-case utility of zero, it constitutes an obviously dominant strategy in this setting.

1.3 Questions Considered in Our Analysis

We ask one question about algorithmic design and three categories of questions about economic design.

1. *How important was it to optimize a solver for the computationally hard problems arising under the specific auction design?* Auctions can include problems that (in general) may not be solvable in a reasonable amount of time; a well-studied example is the winner determination problem in combinatorial auctions [Rothkopf et al., 1998]. Exact solutions can be replaced with approximations, but this may degrade outcomes and incentives. We investigated station repacking, a hard problem that had to be solved repeatedly in the incentive auction. 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? Theoretical arguments sufficed to argue that incentives were robust to failures to solve repacking problems, but the way that such failures would have impacted cost and efficiency was harder to reason about. 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 5.1, showing that substituting the custom feasibility checker with the best off-the-shelf alternative could have increased both average costs and value loss by more than 20%.
2. *How important was it to expand the set of products included in the incentive auction?* In the incentive auction, it would have been straightforward only to buy back licenses from stations broadcasting in the UHF band, since only UHF spectrum needed to be cleared. Instead—considerably increasing complexity—the FCC offered to purchase licenses from stations broadcasting in both the UHF and VHF bands, and offered UHF stations the option of moving into the VHF band rather than going off-air. This increased the pool of stations eligible to participate in the auction by roughly 20%, with the potential both to increase efficiency (a lower-value VHF station could go off-air to make room for a UHF station) and to lower costs (c.f. Bulow and Klemperer's [1996] result that increasing competition can be more important to revenue than setting an optimal reserve price). Was the extra complexity worth it? We found that it was (Section 5.2): under straightforward bidding, not repacking the VHF band could have increased the auction's cost by 20–30% and decreased efficiency by 5–10% (with variation depending both on random sampling and the choice of value model).
3. *Was price discrimination effective at reducing payments to stations and increasing the number of channels cleared?* A canonical insight from revenue maximization [Myerson, 1981] is that bidders with weaker valuation distributions should be boosted to increase competition with stronger bidders. The incentive auction attempted to do something similar, reducing initial price offers for stations that reached smaller populations of viewers. The reduced price offers, called "pops scoring", were politically contentious. In Section 5.3, we show that the effects of scoring were not robust across value models. 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. *What was the impact of the FCC's method of determining supply?* The literature has studied various auction designs in which the seller can adjust supply after observing demand [Back and Zender, 2001]. In the incentive auction, the FCC separately ran forward and reverse auctions at a given clearing target; then, if forward-auction demand was insufficient to cover reverse-auction costs, the clearing target was lowered and a new stage of the auction began. This design was novel and received little advance comment from stakeholders. We compared it to a hypothetical auction having oracle access to the final clearing target. We found that the FCC's clearing procedure led to significantly higher costs and less efficient outcomes (Section 5.4): across all experiments, the clearing procedure increased average value loss by 5–26% and average cost by 4–50% (with variation depending on random sampling, the choice of value model used, and whether the VHF band was repacked). We then introduce a new simple clearing procedure that performs nearly as well without knowing the final clearing target.

The rest of the chapter proceeds as follows. Section 2 describes the Incentive Auction in detail. Section 3 defines our valuation model and bidding model. Section 4 details our experimental setup and Section 5 reports our experimental findings.

2 The Incentive Auction

In this section we describe the Incentive Auction, focusing particularly on the reverse auction. We begin by describing the broad problem addressed by the incentive auction, and then introduce 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 Radio Spectrum Reallocation

We begin by describing the context in which the Incentive Auction arose. Many devices, including broadcast television receivers and cell phones, rely on the transmission of electromagnetic signals. These signals can interfere with each other, so transmission is regulated: e.g., in the US, by the Federal Communications Commission (FCC). Since electromagnetic spectrum suitable for wireless transmission is a scarce resource and since it is difficult for a central authority to assess the relative merits of competing claims on it, since 1994 the FCC has used *spectrum auctions* to allocate broadcast rights (see, e.g., Milgrom [2004]). Many regulators around the world have followed suit. At this point, in the US (as in many other countries), most useful radio spectrum has been allocated. Interest has thus grown in the *reallocation* of radio spectrum from less to more valuable uses. Spectrum currently allocated to broadcast television has received particular attention, for two reasons. First, over-the-air television has been losing popularity with the rise of cable, satellite, and streaming services. Second, the upper UHF frequencies used by TV broadcasters are particularly well suited to wireless data transmission on mobile devices—for which demand is growing rapidly—as they can penetrate walls and travel long distances [Knutson and Gryta, 2014].

It thus made sense for at least some broadcasters to sell their licenses to wireless internet providers willing to pay for them. Ideally, these trades would have occurred bilaterally and without government involvement, as occurs in many other markets. However, two key obstacles made such trade unlikely to produce useful, large-scale spectrum reallocation, both stemming from the fact that wireless internet services require large, contiguous blocks of spectrum to work efficiently. First, a buyer's decision about which block of spectrum to buy would limit the buyer to trading only with broadcasters holding licenses to parts of that block; it could be hard or impossible to find such a block in which all broadcasters were willing to trade. Second, each of these broadcasters would have "holdout power", meaning the broadcaster could demand an exorbitant payment in exchange for allowing the deal to proceed. The likely result would have been very little trade, even if broadcasters valued the spectrum much less than potential buyers.

A 2012 Act of Congress implemented a clever solution to this problem. It guaranteed each broadcaster interference-free coverage in its broadcast area on *some* channel, but not necessarily on its currently used channel. This meant that if a broadcaster was unwilling to sell its license it could instead be moved to another channel, solving the holdout problem. To free up the channel that would permit this move to take place, broadcast rights could be bought from another station in the appropriate geographical area, even if this second station did not hold a license for spectrum due for reallocation. In what follows, we call such an interference-free reassignment of channels to stations a *feasible repacking*.

These trades and channel reassignments were coordinated via a novel spectrum auction run by the FCC between March 2016 and April 2017, dubbed the *incentive auction*. 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.

Roughly, the reverse auction began with a *clearing target* or number of TV channels to decommission. Stations were approached one at a time with a series of decreasing price offers for their broadcast rights. When a station refused an offer, it exited the auction irrevocably and was guaranteed a spot in the leftover channels. As prices fell and more stations declined offers, the leftover channels became more and more congested. Before any station's bid was processed, a "feasibility checker" ensured that the station could still fit in the leftover channels alongside the exited stations without causing undue interference; if it could not, that station was "frozen", meaning that its price stopped falling and it was no longer eligible to exit the auction. The reverse auction further included a provision that allowed stations to exchange their broadcast rights for both a channel in a less desirable (VHF) band and monetary compensation. After the reverse auction concluded, a "forward" ascending clock auction was used to sell licenses in the cleared spectrum to mobile carriers. An outer-loop procedure iterated between reverse and forward auctions to identify the largest possible clearing target for which forward auction revenues exceeded costs.

2.2 The Reverse Auction

We now turn to a fuller and more formal summary of the auction mechanics. Prior to the auction, each television station $s \in S$ in the US and Canada[5] was assigned to a channel $c_s \in \mathcal{C} \subseteq \mathbb{N}$. 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 $\mathrm{pre}(s)$ to refer to a station's pre-auction band, sometimes called a station's home band.

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

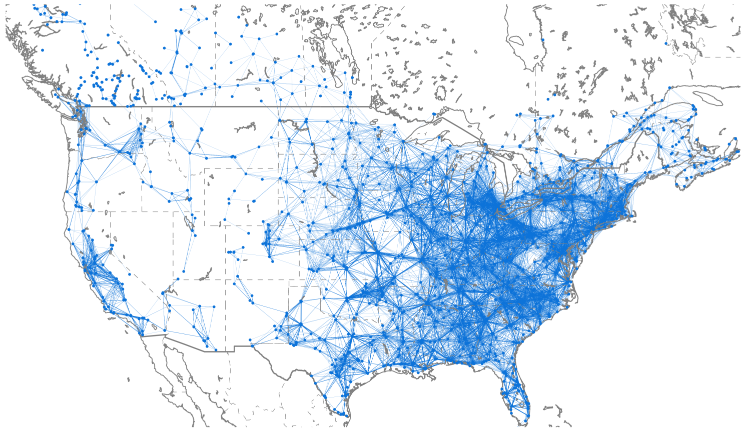

Figure 2: Interference graph derived from the FCC's constraint data.

Each station was 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, 2015a]. 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' could not concurrently be assigned to channels c and c', respectively.

The goal of the reverse 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.

A *feasible assignment* is a mapping $\gamma : \mathcal{S} \to \overline{\mathcal{C}}$ 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.

We now describe a simplified version of the reverse auction in which only the UHF band is repacked: only UHF stations participate in the auction and the only possible 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 things significantly. The complete set of auction rules was published by the FCC in a 230-page document [FCC, 2015b].

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 and (3) processing bids.

A forward auction to sell the cleared spectrum 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 Population (s). We will have more to say about scoring rules in Section 5.3.

The reverse auction is a descending clock auction. At the start of each round, the base clock price is reduced to $p_t = p_{t-1} - d_t$, where $d_t = \max(\frac{p_{t-1}}{20}, \frac{p_0}{100})$. Scores transform base clock prices to individualized station prices; this 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 "winning" station is one that ultimately goes off-air or moves to a different band. 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" $\mathcal{P} : \mathcal{S} \to \mathbb{R}^+$. $\mathcal{P}(s)$ 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

When only the UHF band is repacked, a bid in a given round corresponds to a binary decision. 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". A station that is guaranteed to be frozen for the remainder of the stage is called a *provisional winner*; in a UHF-only auction, every frozen station is provisionally winning. 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

Figure 3: 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.

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 $\sum_{s \in S_{\text{winners}}} \mathcal{P}(s)$),[6] the

[6]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.

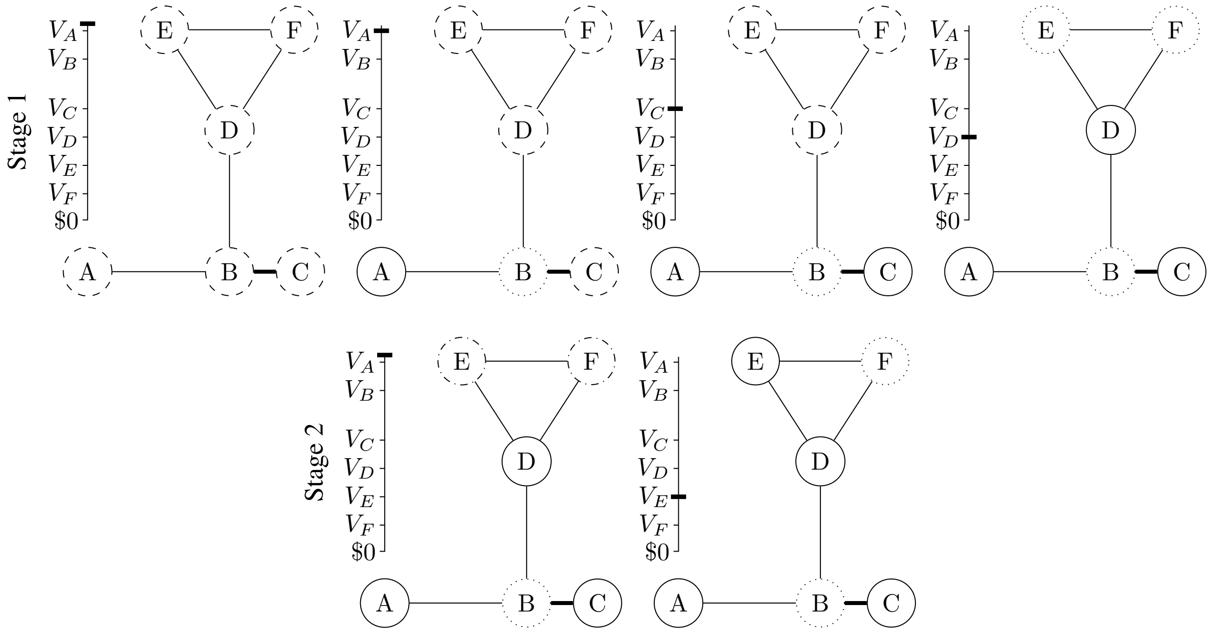

Figure 4: An illustration of the reverse auction example described in Section 2.3. Each row corresponds to a stage; a new image is drawn each time a station exits. Active stations are dashed, exited stations are solid, frozen stations are dotted, and stations with catch-up status are dashdotted. Connected stations cannot jointly broadcast on the same channel. B and C additionally cannot jointly broadcast on adjacent channels (shown by the thick bold edge between them). To the left of each image is the current clock price.

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" 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

Figure 4 illustrates the reverse auction through a worked example. Consider an auction setting with six stations A, B, C, D, E, F having valuations $V_A > V_B > V_C > V_D > V_E > V_F$. Assume that stations bid straightforwardly: they accept offers above their values and reject offers below their values. 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, always finding a feasible assignment when one exists and always determining infeasibility otherwise. 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 each neighboring station according to the interference graph in Figure 4. 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

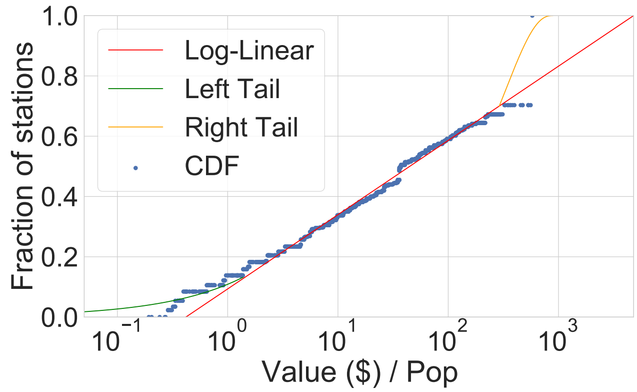

Figure 5: CDF for our maximum likelihood estimate of N and its log-uniform distribution fit, plus generalized Pareto left and right tails.

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 where $\overline{c} = 3$: only one channel is cleared and two channels remain 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 $V_A + V_E$.

3 Value and Bidding Models

This section begins by describing two value models. The first follows Doraszelski et al. [2017]; we created the second for this study, based on 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.

3.1 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., $v_{s,\text{OFF}} = 0$. Both models only provide $v_{s,\text{UHF}}$, 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 $\frac{2}{3} \cdot v_{s,\text{UHF}} \cdot \mathcal{N}(1, 0.05)$ and similarly for the LVHF band as: $\frac{1}{3} \cdot v_{s,\text{UHF}} \cdot \mathcal{N}(1, 0.05)$—i.e., roughly two thirds and one third of the station's UHF value with some multiplicative Gaussian noise. We generated values for VHF stations by computing a hypothetical UHF value and then applying the fractional reductions for VHF bands just described.

3.1.1 The MCS Value Model.

Doraszelski et al.'s [2017] valuation model, which we dub the MCS (Max of Cash flow and Stick value) model, treats a station's value as the maximum of its cash flow value as a business and its stick value. 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. Both of these were estimated from various sources including transaction data of station sales, advertising revenue, and station features.

3.1.2 A Novel Value Model based on Bid Data.

Two years after the incentive auction concluded, the FCC released the auction bids. We used this data to construct a "realistic" model for station valuations. While the bids are not sufficient to reveal station values, they do allow us to infer bounds on values. In some cases these bounds are relatively tight, with upper and lower bounds separated by a single clock interval. Most of the time they are looser, in some cases, to an extent that does not allow us to improve on the trivial bound.

We inferred bounds on each UHF station's home band value, $v_{s,\text{UHF}}$. We then used these bounds to fit a model. We assumed that 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. We computed 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. The results (Figure 5) 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).

To generate UHF values, we multiply a station's population by a sample from the modeled distribution of N. In what follows, we refer to this model as the BD (bid data) model.

3.2 Bidding Model

We now describe our model of how stations bid. A station participated in our simulations if its opening price for going off-air exceeded its value for continuing to broadcast in its home band. After excluding 64 non-mainland stations, we considered 1813 stations eligible to participate in our simulations: 1407 UHF, 367 HVHF, and 39 LVHF.

In auctions with UHF options only, in which bidders are able only to remain off-air or to 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, $\arg\max_{b \in B_{s,t}} P_{s;b;t} + v_{s;b}$. 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.

4 Simulator Design and 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 3. 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. Except when otherwise noted, we started our simulations at the 84 MHz clearing target (the clearing target at which the real incentive auction concluded; see Figure 3 for more details on possible clearing targets) 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 5.4.

The incentive auction's design requires it to begin from a feasible channel assignment for the non-participating stations, however based on which stations participate, this will not always be possible. The FCC's rules therefore allowed a small number of stations to be assigned to channels *within* the spectrum that was otherwise resold (i.e., channels in $\mathcal{C} \setminus \overline{\mathcal{C}}$), even though doing so degrades the desirability of the mobile broadband licenses sold in the forward auction. Such stations are referred to as "impairing". We impair stations as necessary in our simulations to produce a valid starting assignment (for details, see the full paper).

We now explain how we compare simulated outcomes and discuss some simplifications our simulations make relative to the real auction process.

4.1 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 $\sum_{s \in \mathcal{S}} v_{s,\text{pre}(s)} - v_{s,\text{post}(\gamma,s)}$. Ideally, our metric for allocative efficiency would be the ratio of the value loss of a simulation's final assignment, γ, relative to an assignment from an efficient repacking γ^* that minimizes the value loss for a given value profile, i.e., $\frac{\sum_{s \in \mathcal{S}} v_{s,\text{pre}(s)} - v_{s,\text{post}(\gamma,s)}}{\sum_{s \in \mathcal{S}} v_{s,\text{pre}(s)} - v_{s,\text{post}(\gamma^*,s)}}$. In general, however, γ^* is too difficult to compute, so we convert our absolute metric into a relative one by 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 of reaching the specified clearing target: 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 high efficiency (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 Experiments

Our experiments are divided into four categories, based on which element of the auction design they investigate: (1) repacking the VHF band in 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.

5.1 Feasibility Checking

The feasibility checker determines if a given station can be repacked alongside the stations that have previously 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 loss of the allocation is changed.

Do auctions achieve better efficiency and/or lower costs as the feasibility checker improves? Our intuition is that in general, better feasibility checking should lead to better results.

SATFC 2.3.1, the feasibility checker that was used in the incentive auction, was designed over several years [Fréchette et al., 2016, Newman et al., 2017]. The solver is a prime example of AI empowering market design. It combines complete and local-search SAT-encoded feasibility checking with a wide range of domain-specific techniques, such as constraint graph decomposition. Automatic algorithm configuration techniques were applied 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 alternative solvers.

Newman et al. [2017] 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 their reverse auction simulations.[7] The results are shown in Figure 6. We selected certain solvers from this figure and ran simulations, using each as the feasibility checker. Specifically, we compared 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);

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

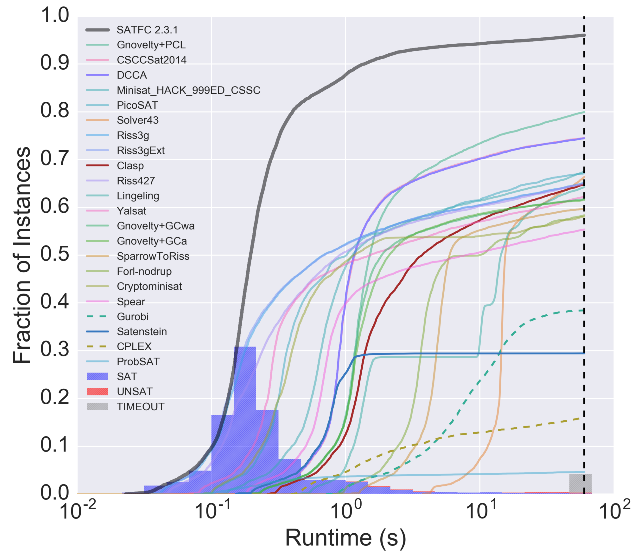

Figure 6: Empirical Cumulative Density Function of runtimes for default configurations of MIP and SAT solvers and for SATFC 2.3.1 on a benchmark set of 10 000 non-trivial problems. Figure taken from Newman et al. [2017]. The curves show fraction of instances solved (y axis) within different amounts of 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.

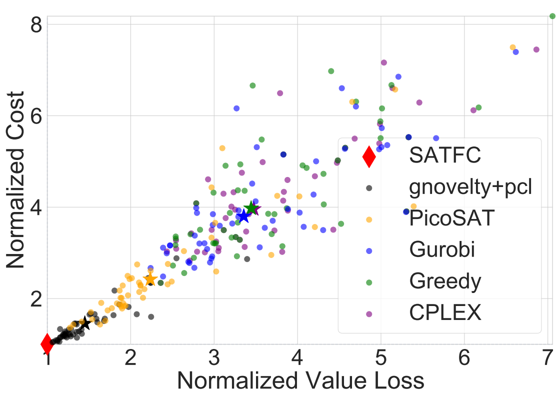

Figure 7: Comparing auctions using different feasibility checkers.

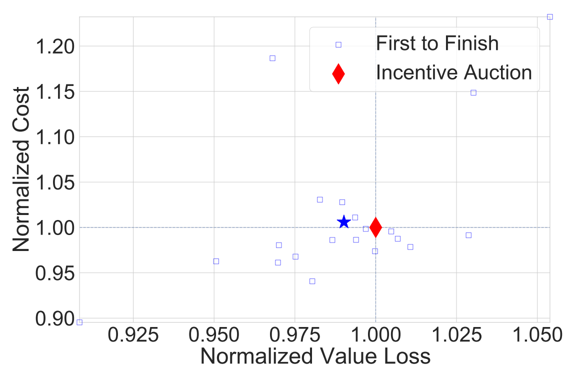

Figure 8: Comparing the first to finish algorithm against the standard bid processing algorithm for single-stage 126 MHz auctions.

- *PicoSAT*: to our knowledge, alongside MIP approaches, 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* and *CPLEX*: MIP solvers initially considered by the FCC; and
- *Gnovelty+PCL*: the best performing of the 22 ACLib solvers on the benchmark data described above.

Our experiments took just over two CPU years. Results are shown in Figure 7. (In what follows, we typically present figures for repacking both UHF and VHF using the BD value model, but discuss our findings across all experimental settings.) Each point in the figure 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 using SATFC 2.3.1) 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

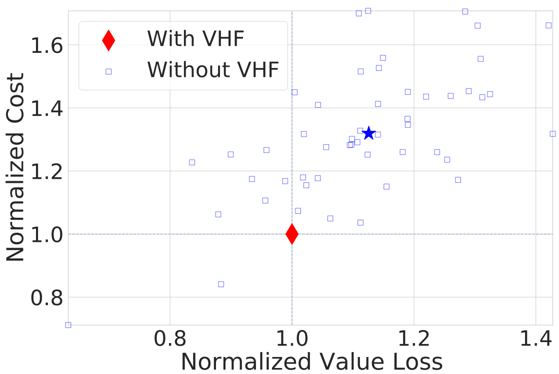

Figure 9: Comparing auctions that only repack the UHF band against auctions that also repack VHF bands.

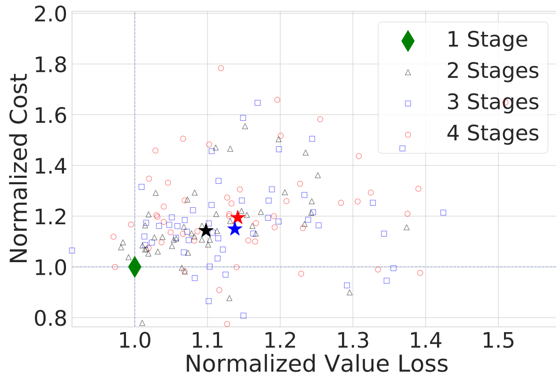

Figure 10: Comparing auctions running through 1-4 stages, ultimately ending on the same clearing target.

treatment we also plot a star to indicate each metric's mean value. We observed that stronger feasibility checkers led to better outcomes according to both of our metrics: the relative rankings of the solvers in the benchmark study translated 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 all other solvers on both metrics, not only on average but in each individual simulation and across both value models. Reverse auctions run using the best off-the-shelf solver, *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.

5.2 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 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. Stations no longer possessed obviously dominant strategies and might have benefited from reasoning about when to move between bands (a "ladder" constraint restricted movement from higher-value to lower-value bands). Price calculations became more involved as each option had to be priced appropriately. The bidding language had to be augmented with "fallback bids" that determine where a station goes if a move to a VHF band is unsuccessful. 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 raised the possibility of more efficient allocations, such gains might have been small, arising only under exotic bidding behavior, or have come at a high cost. We thus asked: 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 9. 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.

5.3 Scoring Rules

Stations' starting prices in the incentive auction were not all the same: they were set proportionally to an assigned *score*, determined by a *scoring rule*. Under truthful bidding in a UHF-only auction, an unfrozen station will remain off-air for exactly as many rounds as it takes for its price to fall below its value, at which point it will exit. Using scoring to raise or lower a station's initial price relative to others is one control knob an auctioneer has over the order in which stations exit.

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 interfere with fewer neighbours. In the FCC's design, this was achieved by offering higher prices to stations with more "interference links" so that those stations would be less likely to exit. Second, scoring reduces an auction's expected cost by offering lower prices to stations that would be likelier to accept them, following Myerson [1981]. In the FCC's design, this was implemented by reducing prices offered to stations serving smaller populations. (This design element was controversial: 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.) Overall, the two elements were combined by setting opening prices in proportion to the square root of the product of a station's population and its interference links.

In our experiments, we compared the following four scoring rules: (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", scoring each station identically. 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 subtlety in these experiments: scoring rules impact prices, so simulations differing only in their scoring rules will not necessarily select the same set of impairing stations given the same value profile. To prevent "apples-to-oranges" comparisons across auctions clearing spectrum of varying quality, we disabled our impairment mechanism and ran simulations with high enough base clock prices to elicit full UHF-station participation. 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 12. 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. 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 interference scoring (97% and 95% of the FCC scoring rule). On the whole, beyond a single setting (MCSMCS values and UHF + VHF repacking) we did not find robust evidence for population-based scoring. These experiments required 3.5 years of CPU time to run.

5.4 Multi-Stage Clearing

The incentive auction allowed for multiple stages of reverse followed by forward auctions in order to let market forces determine the appropriate amount of spectrum to clear. Each successive stage began with a reverse auction clearing a smaller amount of spectrum to achieve a lower overall cost than the previous stage. The actual incentive auction went through four stages.

One obvious practical drawback of a multi-stage approach is its impact on auction length: running multiple stages takes time. A less obvious concern is how the multi-stage approach impacted the final outcome. We 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?

5.4.1 The Economic Impact of Multi-Stage Clearing.

To assess the economic impact of multi-stage clearing, we ran experiments that began 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 3), and each terminating at 84 MHz, 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. As above, we consider a world with high base clock prices to avoid impairments. Under this assumption we observed that running the auction through multiple stages degraded both cost and efficiency, especially when the VHF band was repacked (Figure 10). In VHF-repacking simulations using the MCS (BD) model, on average four-stage auctions cost 1.50 (1.19) times as much as their single-stage counterparts and had 1.26 (1.14) times the value loss. The results were similar but less dramatic in magnitude for UHF-only auctions. In UHF-only simulations using the MCS (BD) model, four-stage auctions 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.

5.4.2 Early Stopping

The results just presented led us to ask: is there any other way the FCC could have determined a clearing target endogenously with less impact on cost and efficiency? We propose a simple answer, which we call *early stopping*: for each candidate clearing target, conduct the forward auction before the corresponding reverse auction, and stop each reverse auction as soon as its cost exceeds the forward auction revenue. To see why this would help, consider again our example in Figure 4, but now 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 exceeds the forward auction revenue. Following our proposal, the first reverse auction 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 second stage target. In general, continuing the reverse auction instead of aborting can only result in more stations exiting, limiting flexibility in later stages. The original design proposed that the forward and reverse auctions be run in parallel, but due to finite staffing resources, the auctions had to be sequenced. We note that early stopping represents a potential algorithmic improvement even over the original parallel design.

Early stopping does not always outperform the original design; it is possible to construct examples where the commitments made in earlier stages are better than those made in later stages. Nevertheless, we believed 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. The second set 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 inputs. We had little data to use for modeling these revenues—one observation for each of the four clearing targets that were reached in practice—so we adopted a convenient regression model. Unlike previous experiments where we assumed a predetermined number of stages, we now determined the auction’s end by comparing forward auction revenues to clearing costs. This meant that auctions could end at any stage, so the amount of spectrum sold could differ across paired simulations. We assume that clearing more spectrum is preferable to clearing less, noting the FCC’s stated goals. We ran 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 3) from that point 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 11. For computational reasons, we only ran UHF-only simulations; even so, the experiments took more than 15 CPU years overall. To avoid making comparisons between different qualities of spectrum, we disabled our impairment mechanism and let prices start as high as required for full participation. We observed that in both value models, early stopping performed well relative to 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 clearing procedure to determine which cleared more spectrum. Under both value models, on average more mobile licenses were created when by early stopping than the original algorithm—0.50 and 0.22 extra licenses under the BD and MCS models, respectively. Early stopping also led to shorter auctions, averaging 30% and 76% of the rounds required by the original algorithm under BD and MCS, respectively.

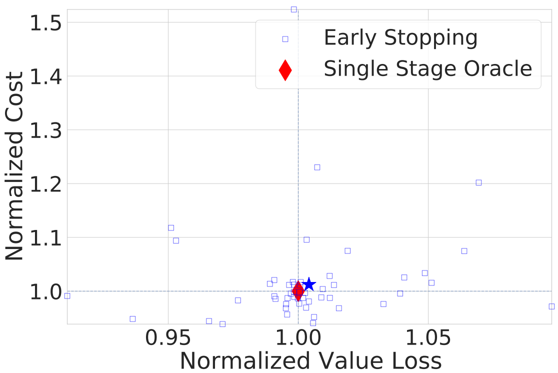

Figure 11: Comparing early stopping auctions against single-stage auctions.

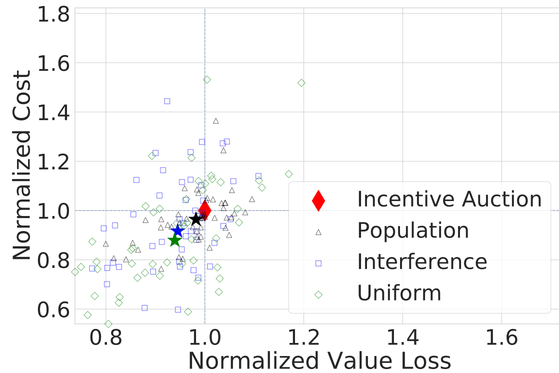

Figure 12: Comparing auctions using four different scoring rules.

We conclude by reflecting on some concrete numbers from the real auction. At the end of the 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.

6 Conclusions and Reflections on the Intersection Between AI and Market Design

High fidelity, computationally intensive simulations are an approach for quantitatively contrasting candidate market designs which offer a key advantage: the quality of answers provided by such simulations will improve with both available computational resources and available market data. We outlined a six-step simulation methodology and instantiated it in the context of the incentive auction to investigate previously unanswerable questions about the cost and efficiency of certain alternative designs. To validate the robustness of our results, we used two quite different value models: one from the empirical economics literature and another that we constructed to rationalize public bid data. Our main findings were that: the specialized feasibility checker developed for the auction significantly improved both cost and efficiency; repacking VHF led to significantly lower costs and more efficient outcomes; 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 multiple stage clearing rule substantially both increased costs and reduced efficiency; and a simple amendment to the clearing algorithm could nearly eliminate multi-round inefficiency. We hope these specific insights can help to inform future auction designs. More broadly, our analysis demonstrates the practicality[8] and promise of large-scale computational analysis of the simulated behavior of candidate market designs in highly complex settings.

Of course, simulations are not the only way that AI can be applied to market design. We offered a second example in passing when discussing SATFC: this algorithm was automatically designed using algorithm configuration and algorithm portfolio techniques, achieving state-of-the-art performance by leveraging data characterizing realistic problem inputs. Some additional ways in which AI can impact the design of practical markets include helping to overcome search frictions problems when matching with an interested counterparty is hard, and helping participants express complex or nonlinear preferences.

The recent success and wide adoption of large language models (LLMs) is leading many to dream bigger. The fundamental catch is counterfactuals: most machine learning techniques find patterns in independent and identically distributed samples, whereas most market design aims to optimize over possible worlds for which no prior data exists. We nevertheless foresee various ways in which machine learning approaches in general, and LLMs in particular, could be useful for market design. The first is by using machine learning to improve individual components of a market-design pipeline. We have already discussed how econometric methods can be used to fit valuation distributions to data. Similarly, observational data could be used to fit models of bidding behavior (e.g., drawing on approaches from

[8]Altogether, our final experiments took more than 60 CPU years to run. Writing this chapter 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.

behavioral game theory [e.g., Camerer, 2003, Wright and Leyton-Brown, 2017]). Second, machine learning could also be used to improve individual components of a market itself. LLMs could be used to elicit preferences, to conduct rich automated negotiations [e.g., Bakhtin et al., 2022], and even to serve as user interfaces on top of messy markets in which bidders might hold complex preferences (e.g., eBay, ad markets). Third, machine learning can be used in a principled way to directly attack the problem of reasoning about counterfactuals. Causal inference has been a hot area for over a decade [e.g., Bottou et al., 2013, Wager and Athey, 2018] . Reinforcement learning simulations of economic systems are becoming rich enough to answer market design questions in some cases [e.g., Zheng et al., 2022]. LLMs even offer the prospect of serving as a cheap proxy for human subject experiments [Aher et al., 2023, Horton, 2023]. Fourth, and finally, AI is likely to give rise to a new application domain for market designers. Important open questions include how to monetize LLMs (via, e.g., product placement or sponsored recommendations); how to defend LLMs against strategic manipulation (e.g., non-sponsored recommendations or prompt injection attacks); how to structure partnerships between custom LLMs and broader algorithmic services within which they are embedded; and how to price access to LLMs (particularly given the existence of a complex quality/cost tradeoff).

References

David Delacrétaz, Scott Duke Kominers, and Alexander Teytelboym. Matching Mechanisms for Refugee Resettlement. Working Papers 2019-078, Human Capital and Economic Opportunity Working Group, December 2019.

Maxwell Allman, Itai Ashlagi, Irene Lo, Juliette Love, Katherine Mentzer, Lulabel Ruiz-Setz, and Henry O'Connell. Designing school choice for diversity in the San Francisco unified school district. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, page 290–291, 2022.

Daniel Freund, Shane G. Henderson, and David B. Shmoys. Minimizing multimodular functions and allocating capacity in bike-sharing systems. *Production and Operations Management*, 27(12):2346–2349, 2018.

Rong Yang, Benjamin Ford, Milind Tambe, and Andrew Lemieux. Adaptive resource allocation for wildlife protection against illegal poachers. In *International Conference on Autonomous Agents and Multi-Agent Systems*, page 453–460, 2014.

Martin Bichler, Vladimir Fux, and Jacob K Goeree. Designing combinatorial exchanges for the reallocation of resource rights. *Proceedings of the National Academy of Sciences*, 116(3):786–791, 2019.

Nicolau Santos, Paolo Tubertini, Ana Viana, and Joao Pedro Pedroso. Kidney exchange simulation and optimization. *Journal of the Operational Research Society*, 68(12):1521–1532, 2017.

Kevin Leyton-Brown, Paul Milgrom, and Ilya Segal. Economics and computer science of a radio spectrum reallocation. *Proceedings of the National Academy of Sciences*, 114(28):7202–7209, 2017. ISSN 0027-8424.

Frank P. Kelly and Richard Steinberg. A combinatorial auction with multiple winners for universal service. *Management Science*, 46:586–596, 2000.

Paul Milgrom and Ilya Segal. Clock auctions and radio spectrum reallocation. *Journal of Political Economy*, 128(1): 1–31, 2020.

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

Sushil Bikhchandani, Sven De Vries, James Schummer, and Rakesh V Vohra. An ascending Vickrey auction for selling bases of a matroid. *Operations research*, 59(2):400–413, 2011.

John A Ferejohn, Robert Forsythe, and Roger G Noll. An experimental analysis of decision making procedures for discrete public goods: A case study of a problem in institutional design. *V. L. Smith, ed. Research in Experimental Economics*, pages 1–58, 1979.

S. J. Rassenti, V. L. Smith, and R. L. Bulfin. A combinatorial auction mechanism for airport time slot allocation. *The Bell Journal of Economics*, 13(2):402–417, 1982. ISSN 0361915X.

Anthony M. Kwasnica, John O. Ledyard, David Porter, and Christine DeMartini. A new and improved design for multiobject iterative auctions. *Management Science*, 51:419–434, 2005.

Timothy N. Cason, Karthik N. Kannan, and Ralph B. Siebert. An experimental study of information revelation policies in sequential auctions. *Management Science*, 57:667–688, 2011.

Gediminas Adomavicius, Shawn P. Curley, Alok Gupta, and Pallab Sanyal. Effect of information feedback on bidder behavior in continuous combinatorial auctions. *Management Science*, 58:811–830, 2012.

Nikhil Agarwal and Paulo Somaini. Demand analysis using strategic reports: An application to a school choice mechanism. *Econometrica*, 86(2):391–444, 2018.

Caterina Calsamiglia, Chao Fu, and Maia Güell. Structural estimation of a model of school choices: The Boston mechanism versus its alternatives. *Journal of Political Economy*, 128(2):642–680, 2020.

Elena Krasnokutskaya and Katja Seim. Bid preference programs and participation in highway procurement auctions. *American Economic Review*, 101(6):2653–2686, 2011.

Paulo Somaini. Identification in auction models with interdependent costs. *Journal of Political Economy*, 128(10): 3820–3871, 2020.

Susan Athey, Jonathan Levin, and Enrique Seira. Comparing open and sealed bid auctions: Evidence from timber auctions. *The Quarterly Journal of Economics*, 126(1):207–257, 2011.

Susan Athey and Denis Nekipelov. A structural model of sponsored search advertising auctions. In *6th Ad Auctions Workshop*, volume 15, 2010.

Jean L Kiddoo, Evan Kwerel, Sasha Javid, Melissa Dunford, Gary M Epstein, Charles E Meisch Jr, Karla L Hoffman, Brian B Smith, Anthony B Coudert, Rudy K Sultana, et al. Operations research enables auction to repurpose television spectrum for next-generation wireless technologies. *INFORMS Journal on Applied Analytics*, 49(1):7–22, 2019.

Michael Kearns and Lili Dworkin. A Computational Study of Feasible Repackings in the FCC Incentive Auctions. *arXiv:1406.4837 [cs]*, June 2014. arXiv: 1406.4837.

Peter Cramton, Hector López, David Malec, and Pacharasut Sujarittanonta. Design of the reverse auction in the FCC Incentive Auction. 2015.

Ulrich Doraszelski, Katja Seim, Michael Sinkinson, and Peichun Wang. Ownership concentration and strategic supply reduction. Technical Report 23034, National Bureau of Economic Research, January 2017.

Lawrence M Ausubel, Christina Aperjis, and Oleg Baranov. Market design and the FCC Incentive Auction. Working Paper, 2017.

FCC. In the matter of expanding the economic and innovation opportunities of spectrum through incentive auctions. *Notice of Proposed Rulemaking*, FCC 12–118, 10 2012.

Michael H. Rothkopf, Aleksandar Sasa Pekec, and Ronald M. Harstad. Computationally manageable combinational auctions. *Management Science*, 44:1131–1147, 1998.

Jeremy Bulow and Paul Klemperer. Auctions versus negotiations. *The American Economic Review*, 86(1):180–194, 1996. ISSN 00028282.

Roger B Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

Kerry Back and Jaime F. Zender. Auctions of divisible goods with endogenous supply. *Economics Letters*, 73(1):29–34, 2001. ISSN 0165-1765.

P. Milgrom. *Putting Auction Theory to Work*. Churchill Lectures in Economics. Cambridge University Press, 2004. ISBN 9781139449168.

Ryan Knutson and Thomas Gryta. Verizon, AT&T May Face Bidding Limits in Spectrum Auction. *Wall Street Journal*, April 2014. ISSN 0099-9660. URL `http://www.wsj.com/articles/SB10001424052702304626304579510154106120342`. `http://www.wsj.com/articles/SB10001424052702304626304579510154106120342`, Accessed 2016-12-12.

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

FCC. Repacking constraint files. `http://data.FCC.gov/download/incentive-auctions/Constraint_Files/`, 2015a. Accessed 2015-11-20.

FCC. Application procedures for broadcast Incentive Auction scheduled to begin on March 29, 2016. *FCC Public Notice*, DA 15-1183, 10 2015b.

Alexandre Fréchette, Neil Newman, and Kevin Leyton-Brown. Solving the station repacking problem. In *AAAI Conference on Artificial Intelligence*, pages 702–709, 2016.

Neil Newman, Alexandre Fréchette, and Kevin Leyton-Brown. Deep optimization for spectrum repacking. *Communications of the ACM*, 61(1):97–104, 2017.

Frank Hutter, Manuel López-Ibáñez, Chris Fawcett, Marius Lindauer, Holger H Hoos, Kevin Leyton-Brown, and Thomas Stützle. AClib: A benchmark library for algorithm configuration. In *Conference on Learning and Intelligent Optimization (LION)*, pages 36–40. Springer, 2014.

Kevin Leyton-Brown. The viability of exact feasibility checking. 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, February 2013. URL `https://www.cs.ubc.ca/~kevinlb/talks/2013-Feasibility-Testing.pdf`.

Colin F. Camerer. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press, 2003.

James R Wright and Kevin Leyton-Brown. Predicting human behavior in unrepeated, simultaneous-move games. *Games and Economic Behavior*, 106:16–37, 2017.

Anton Bakhtin, Noam Brown, Emily Dinan, Gabriele Farina, Colin Flaherty, Daniel Fried, Andrew Goff, Jonathan Gray, Hengyuan Hu, et al. Human-level play in the game of diplomacy by combining language models with strategic reasoning. *Science*, 378(6624):1067–1074, 2022.

Léon Bottou, Jonas Peters, Joaquin Quiñonero-Candela, Denis X. Charles, D. Max Chickering, Elon Portugaly, Dipankar Ray, Patrice Simard, and Ed Snelson. Counterfactual reasoning and learning systems: The example of computational advertising. *Journal of Machine Learning Research*, 14(101):3207–3260, 2013. URL `http://jmlr.org/papers/v14/bottou13a.html`.

Stefan Wager and Susan Athey. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242, 2018.

Stephan Zheng, Alexander Trott, Sunil Srinivasa, David C Parkes, and Richard Socher. The ai economist: Taxation policy design via two-level deep multiagent reinforcement learning. *Science advances*, 8(18):eabk2607, 2022.

Gati Aher, Rosa I. Arriaga, and Adam Tauman Kalai. Using large language models to simulate multiple humans and replicate human subject studies, 2023.

John J. Horton. Large language models as simulated economic agents: What can we learn from homo silicus?, 2023.