ARTIFICIAL INTELLIGENCE AND MARKET DESIGN

PAUL MILGROM & STEVEN TADELIS

SEPTEMBER 2017

OUTLINE OF PAPER

1. Outline
2. Machine learning and the Incentive Auction
   - Auctions for complex resource allocation
   - Role of ML in the US incentive auction
3. Online bidding for programmatic advertising
   - Using ML to learn valuations
   - Using ML to learn about competitors
4. Bidding with No Regret Strategies
5. Managed Marketplaces
   - Using ML/NLP for quality control
   - Using NLP for feedback evaluation
   - Personalized Search
6. Other Topics
   - Shopbots for price comparisons
   - Pricebots for price discrimination
   - Recommendations (complements, airline tickets)

COMPUTER SCIENCE & THE "INCENTIVE AUCTION"

…and AI/ML, of course!

4 INCENTIVE AUCTION BACKGROUND

- US National broadband plan (2010)
  - Growing demand for radio spectrum capacity for broadband and related services.
  - Shrinking value of licenses used for over-the-air TV broadcasts.
  - *Hypothesis: It is efficient to reallocate some channels to broadband.*
- Spectrum reallocation is a collective action problem with a role for government
  - International and inter-regional coordination of frequency uses.
  - "Zoning-like" restrictions on adjacent uses.
  - Band plan (uplink/downlink/guard bands) depends on bandwidth reallocated.
  - All-at-once timing of channel switches by non-sellers.

5 ENDED IN APRIL 2017

OUTCOME OF THE INCENTIVE AUCTION

- Lots of Spectrum
  - Cleared fourteen TV channels 38-51 (84MHz)
  - 70MHz for use in mobile broadband
  - 14MHz for unlicensed uses
- Lots of Money
  - Gross auction revenue of $19.8 billion
  - 175 winning broadcasters received $10.05 billion
  - Highest price was $304 million
  - 11 non-commercial stations received more than $100 million
  - KQED (SF public television) received $95 million

6

CO-CHANNEL INTERFERENCE CONSTRAINTS

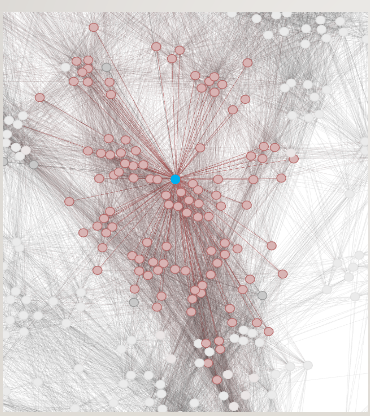

Each node is a UHF-TV station.

Each arc is a pair of stations that cannot be both assigned to the same channel.

Nodes connected to the *blue node* are colored in pink.

7

FEASIBILITY CHECKING IS NP-COMPLETE

CO-CHANNEL INTERFERENCE IN THE US AND CANADA

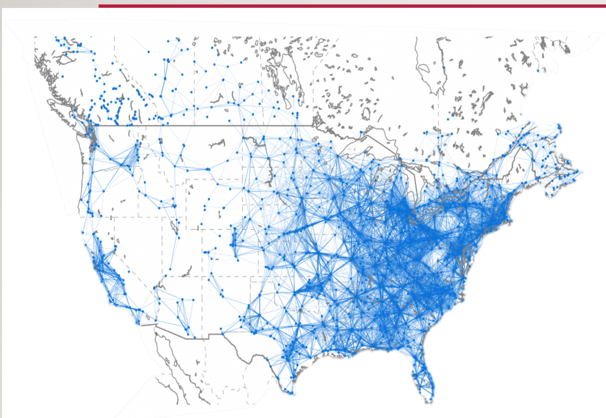

- Question: Using only channels 14-36, is it possible to assign channels to stations in set S without encountering interference?
  - About 75,000 such questions in the auction.
- There are about 130,000 "co-channel constraints" shown in the graph.
  - Graph coloring is an NP-complete problem.
- Actual constraints are more detailed and numerous.
  - About 2.7 million constraints in the full list.

8

OPTIMIZATION IS "HARDER" AND LIMITS POSSIBLE AUCTION DESIGNS!

VICKREY PRICE COMPUTATIONS

- Let $S \in \mathcal{F}$ mean that S is a feasible set of broadcasters.
- Then, the Vickrey price for a station i that goes off air is

$$p_i = \left(\max_{S\in\mathcal{F}} \sum\nolimits_{j\in S} v_j\right) - \left(\max_{\substack{S\in\mathcal{F}\\ S\ni i}} \sum\nolimits_{j\in S} v_j\right)$$

- With 2000 stations, a 1% computation error in one of the maximizations leads to a pricing error of $\approx 20\times$average station value.
- *$\therefore$Vickrey prices are not computable in practice.*

9 COMPLEXITY AND THE HUMAN INTERFACE

Dear Mr. Broadcaster:

We have heard your concerns about the complexity of the spectrum reallocation process. You may even be unsure about whether to participate or how much to bid. To make things as easy as possible for you, we have adopted a Nobel-prize winning auction procedure called the "Vickrey auction."

In this auction, all you need to do is to tell us what your broadcast rights are worth to you. We'll figure out whether you are a winner and, if so, how much to pay to buy your rights. The rules will ensure that it is in your interest to report truthfully. That is the magic of the Vickrey auction!

The computations that we do will be very hard ones, and we cannot guarantee that they will be exactly correct. Also, federal law forbids us to share the information that you would need to check them.

…. [*Read about the alternative auction design in Leyton-Brown, Milgrom & Segal (2017). But, we still need feasibility checking!*]

KLB's lab

WORK AND SLIDES BY KEVIN LEYTON-BROWN'S LAB

FEASIBILITY CHECKING: EXPERIMENTS

- **Data** generated Nov 2015 – Feb 2016 using
  - the FCC's Nov 2015 interference constraints
  - an auction simulator
  - varying **simulation assumptions**:
    - how much spectrum is cleared: 126 MHz; 108 MHz; 84 MHz
    - which stations opt to participate
    - these stations' valuations
    - the timeout given to SATFC in the simulation (1; 5; 10; 60 min)
- 128 **auctions** ⇒ 1.4 M **instances**
  - 6,128 – 17,764 instances per auction
    - about 80% solvable by directly augmenting the previous solution
    - Focus on the other 20%
  - split auctions 102/26 into training/test sets
- Our goal: solve problems within a **one-minute cutoff**

KLB's lab

OPTIMIZATION (MIP) VS SAT APPROACHES

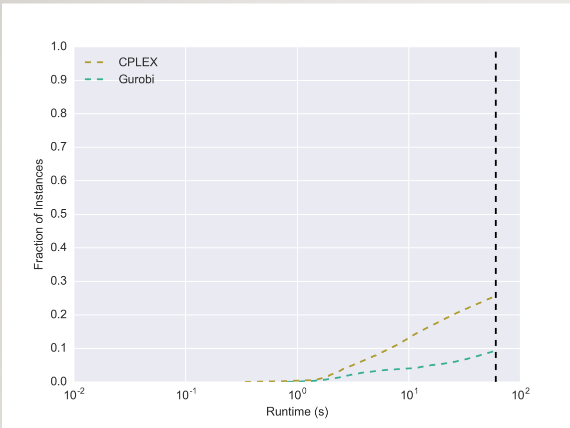

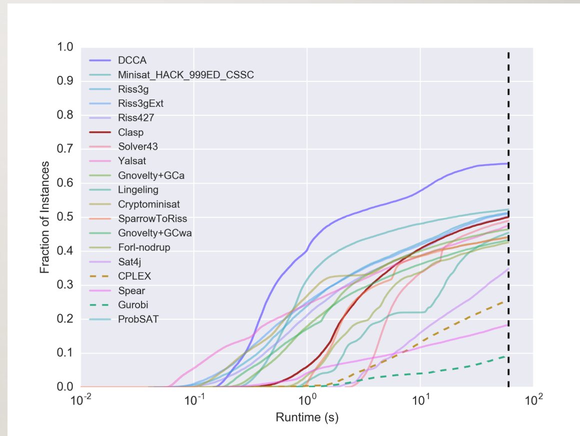

KLB's lab

MACHINE LEARNING

ALGORITHM CONFIGURATOR

- "Deep optimization" uses automated methods to choose algorithm designs from a **highly parameterized space**
  - which branching heuristic, variable ordering, preprocessing strategy, clause learning technique,

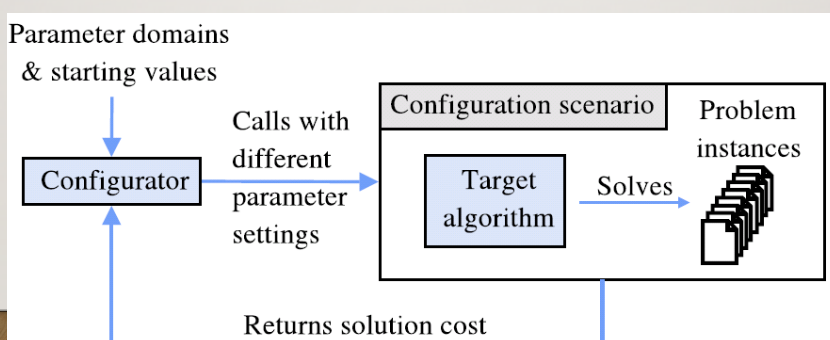

KLB's lab

SEQUENTIAL MODEL-BASED ALGORITHM CONFIGUARATION (SMAC)

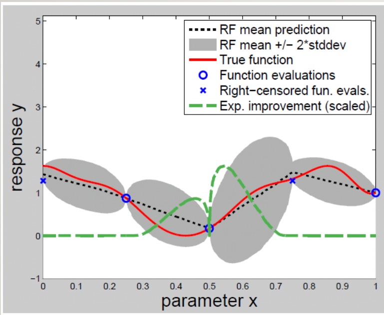

- Initialize with a single run for the default configuration
- Repeat until time budget is exhausted
  1. Learn a random forest model $m: \Theta \times \Pi \rightarrow \mathbb{R}$ from data so far
  2. Compute mean performance $f(\theta) = E_\pi[m(\theta, \pi)]$.
  3. Find θ^* that maximizes *expected improvement*
  4. Add new data using θ^*

14 BUILDING SATFC (SOFTWARE)

[Auctionomics team: Kevin Leyton-Brown, with Alex Frechette and Neil Newman]

1. Pre-identifying "unconstrained stations" to decompose the interference graph into smaller graphs to be solved separately.
2. Training a parameterized heuristic ("CLASP") using machine learning to run fast on instances generated by simulations.
3. Creating a portfolio of algorithms with run times that time-out on "largely disjoint" sets of of instances.
4. Including a "local solver" in the portfolio
5. Including in the portfolio a searchable cache of problem components with known solutions.

KLB's lab

SATFC'S PERFORMANCE

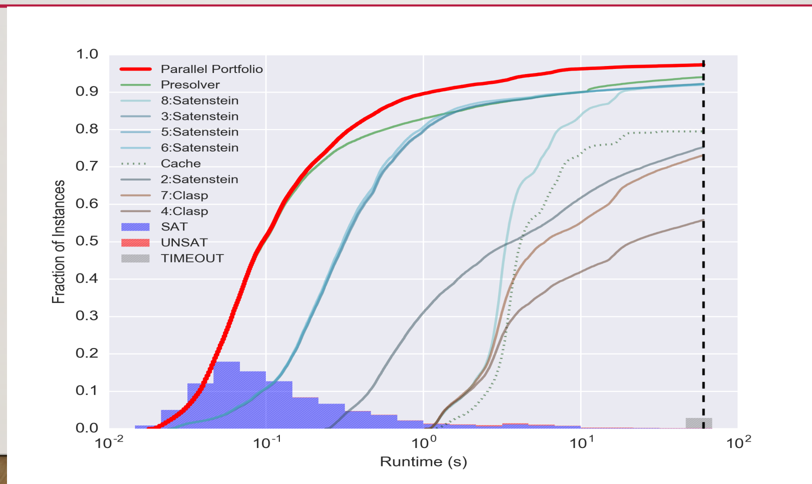

Milgrom, Segal & KLB

AUCTION PERFORMANCE

- We cannot compute VCG on the full problem, so…
  - We restrict attention to the 218 stations within **two links of New York City**
    - a very densely connected region
- Reverse auction simulator (UHF only)
- Simulation **assumptions**:
  - 100% participation
  - 126 MHz clearing target
  - valuations from by a prominent stochastic model
  - 1 min timeout given to SATFC

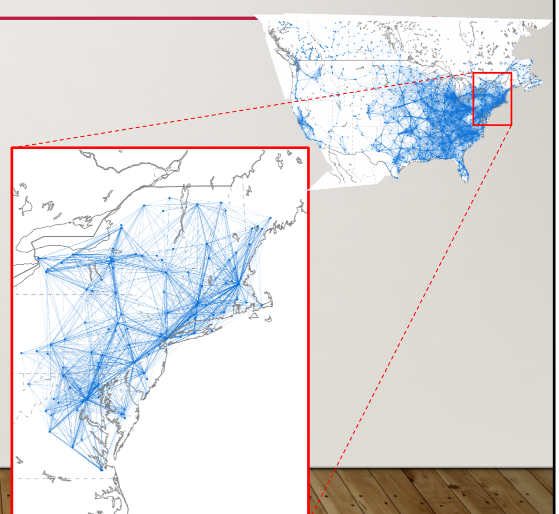

17 *COMPUTATIONAL* PERFORMANCE

- Computation times using Intel Xeon R Processors E5-2640 v2.
  - Optimum and Vickrey prices: 108 hours.
  - Clock solution and prices: 7 minutes.

EFFICIENCY & *COST* PERFORMANCE

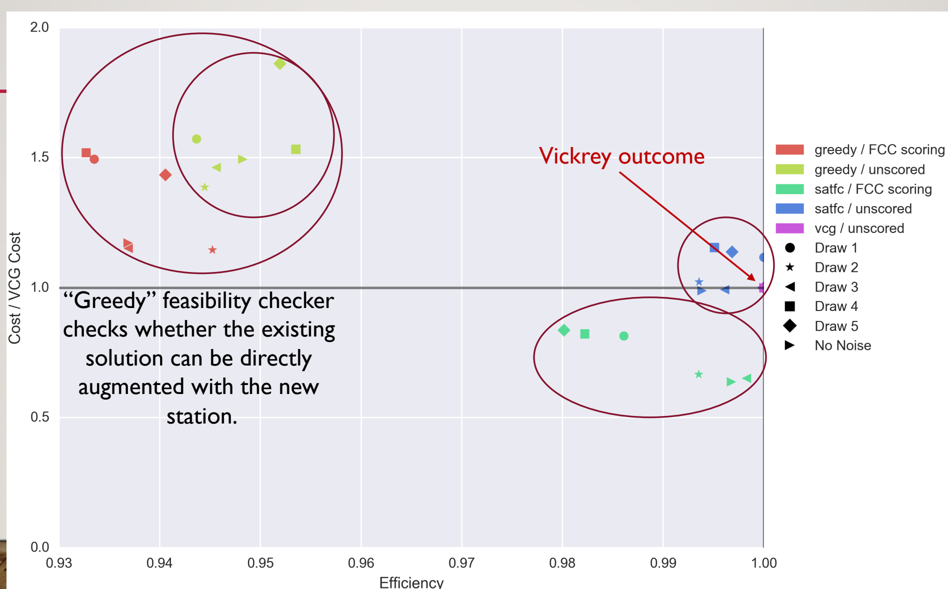

19 SPECIAL THANKS

- Evan Kwerel (FCC). Conceived the incentive auction and the property rights that resolve the hold-out problem.
- Ilya Segal (Stanford). Conceived the general deferred acceptance auction and many of its extensions.
- Kevin Leyton-Brown (UBC). Developed fast solutions that solved ~75,000 NP-hard problems at a enormous scale.

ML/NLP AND MANAGED MARKETPLACES

…and AI/ML, of course!

21 ML/NLP IN MANAGED MARKETPLACES: A. MEASURING SELLER QUALITY

- Feedback systems are credited with being the lubricant of trust in online marketplaces, but "grade inflation" is rampant
- How can a marketplace use behavioral data to measure the quality of a transaction?
- When buyers and sellers interact with messages, there is a treasure trove of post-transaction data
- NLP/Semantic analysis can be used to tease out buyer satisfaction and other measures of a transaction's quality

22 SENTIMENT AND SATISFACTION ON EBAY

"Canary in the e-Commerce Coal Mine: Detecting and Predicting Poor Experiences Using Buyer-to-Seller Messages" (with Dimitriy V. Masterov and Uwe F. Mayer)16th ACM Conference on Electronic Commerce, (EC 2015), June 2015

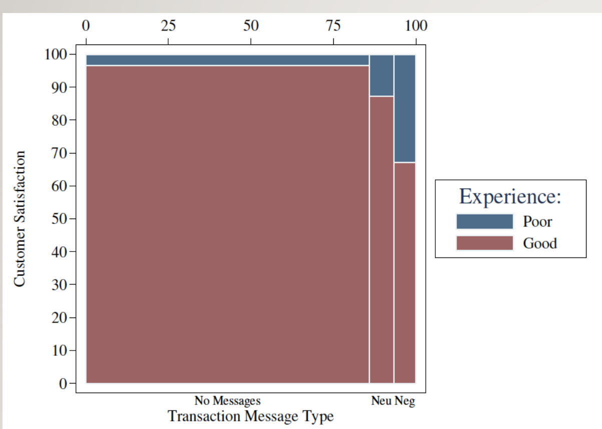

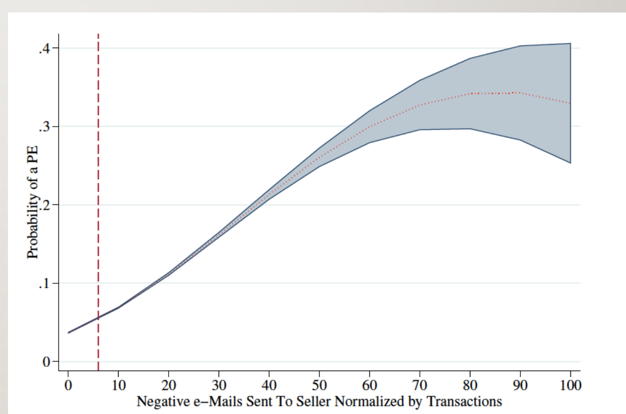

Sample mean shown by vertical dashed line. Graph excludes sellers with fewer than 5 transactions in the previous year.

23 ML/NLP IN MANAGED MARKETPLACES : B. CREATING A MARKET FOR FEEDBACK

- Informative feedback is valuable but costly to provide
- What if a marketplace allows sellers to pay for feedback?
- The marketplace has to mediate the market for feedback to get truthful and informative feedback
- NLP/content analysis can be used to measure the quality and relevance of feedback provided
- Taobao did exactly this to manage their "Reward For Feedback" mechanism

24 TAOBAO'S REWARDS FOR FEEDBACK

Buying Reputation as a Signal of Quality: Evidence from an Online Marketplace (with Lingfang (Ivy) Li and Xiaolan Zhou)

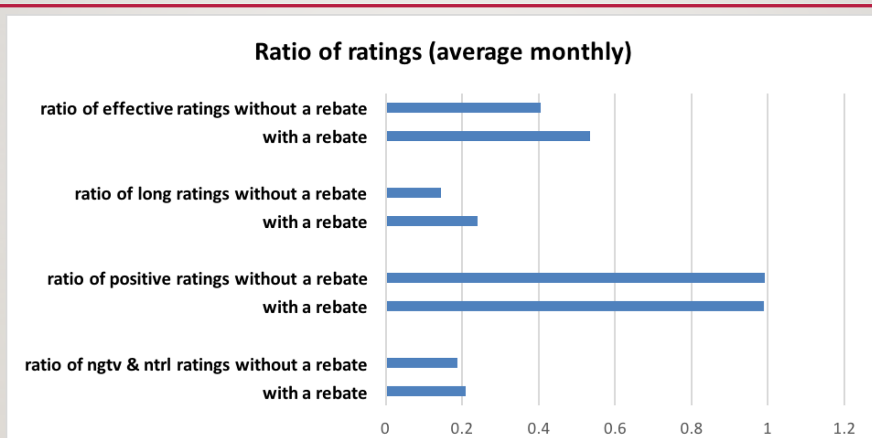

25 ML/NLP IN MANAGED MARKETPLACES : C. PERSONALIZED SEARCH

- Search is a either an exploratory experience, or a friction en route to purchase
- What if a marketplace can figure out tastes and segment consumers?
- Use ML to classify consumer segments and infer best ”consideration sets” for exploratory search
- Orbitz: Mac vs. PC users see different sets (execs quoted saying “won’t show the exact same room two different customers at different prices”)

26 FURTHER TOPICS

- Good recommendation systems can help build customer loyalty
- Retailers/marketplaces may engage in ML-based price discrimination
- But consumers may have access to crowd-sourced shop-bots
- And there are the unknown unknowns…

27

THE END