Adaptive Mechanism Design: A Metalearning Approach

David Pardoe[1]

Peter Stone[1]

Maytal Saar-Tsechansky[2]

Kerem Tomak[2]

[1]Department of Computer Sciences

[2]McCombs School of Business

The University of Texas at Austin

Auction Example

- Consider a book seller using an auction service
- Seller must choose parameters defining auction
- Goal is to maximize revenue
- Optimal parameters depend on bidder population

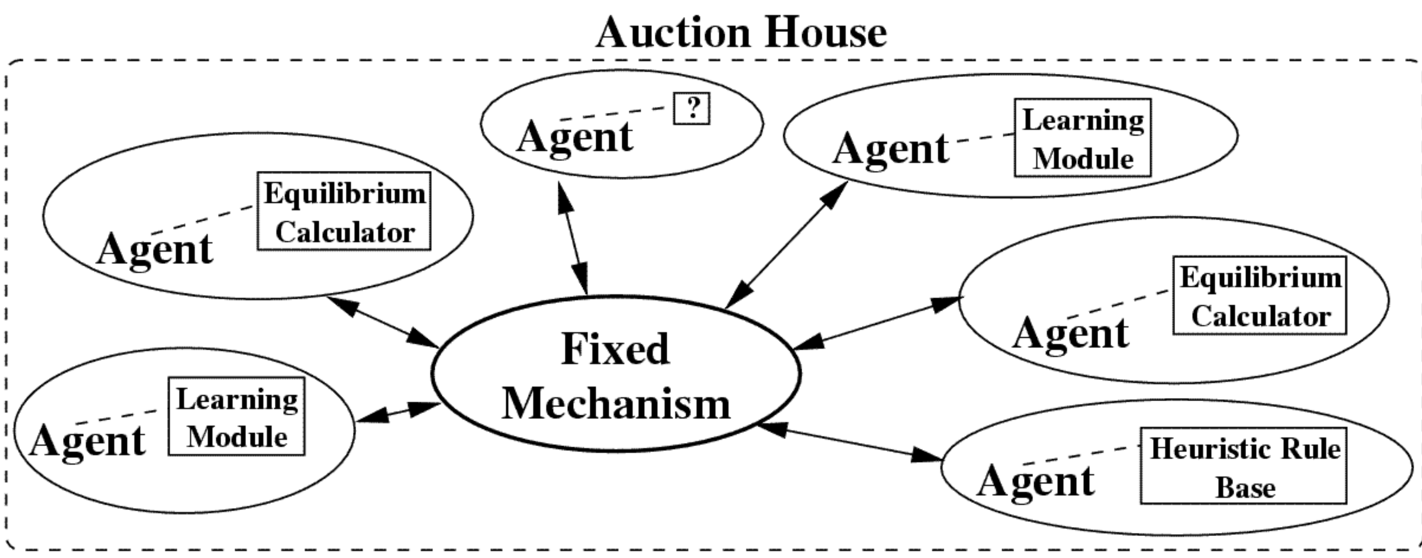

Analytical Approach

- Traditional approach
  - (e.g. Myerson 81, Milgrom and Weber 82)
- Assumptions are made about
  - bidder motivations (valuations, risk aversion, etc.)
  - information available to bidders
  - bidder rationality
- Derive equilibrium strategies
- What if assumptions are incorrect?
  - revise assumptions
  - requires time and human input
  - problem if limited time between auctions

Empirical Approach

- Possible if historical data on similar auctions
- Do data mining to identify optimal parameters (e.g. Shmueli 05)
  - a number of businesses provide this service

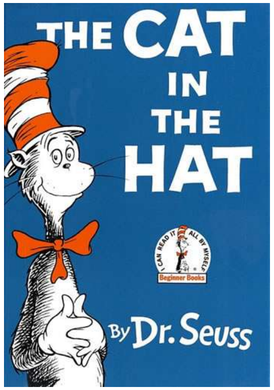

For "The Cat in the Hat", you should run

a 3-day auction starting on Thursday

with a starting bid of $5.

Empirical Approach

- What if the item is new and no data exists?

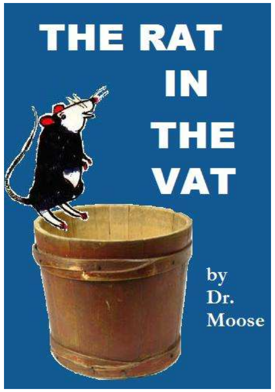

- What if there is a sudden change in demand?

Overview

- Motivation
- Adaptive auction mechanisms
- Bidding scenario
- Adaptive mechanism implementation and results
- Incorporating predictions through metalearning
- Additional experiments

Adaptive Auction Mechanisms

- For use in situations with recurring auctions
  - repeated eBay auctions, Google keyword auctions, etc.
- Bidder behavior consistent for some period
  - possible to learn about behavior through experience
- Adapt mechanism parameters in response to auction outcomes in order to maximize some objective function (such as seller revenue)

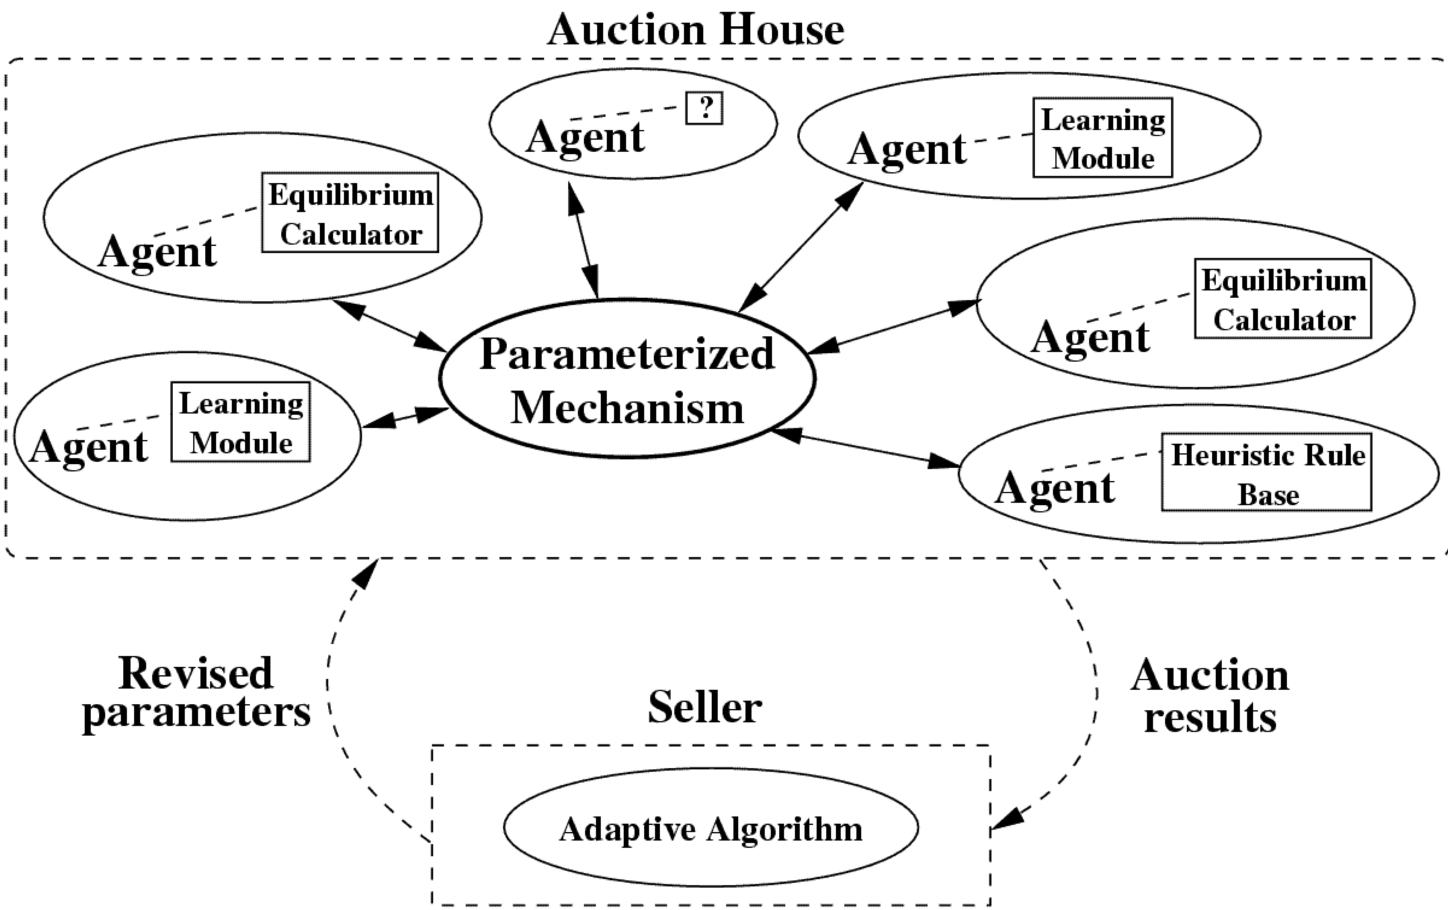

- Seller adjusts parameters using an adaptive algorithm
  - characterizes function from parameters to results
  - essentially an active, online regression learner

Adaptive Auction Mechanisms

- Related work (e.g. Blum et al. 03)
  - apply online learning methods
  - few or no assumptions about bidders
  - worst case bounds
- What about the intermediate case?
  - between complete knowledge and no knowledge
  - can make some predictions about bidders
  - choose adaptive algorithm using this information

Overview

- Motivation
- Adaptive auction mechanisms
- Bidding scenario
- Adaptive mechanism implementation and results
- Incorporating predictions through metalearning
- Additional experiments

Loss Averse Bidders

- Loss aversion: utility of gain = X,
  utility of loss = - X
- Loss averse bidders “lose” if outbid after they were the high bidder
- 2 bidder equilibrium: (Dodonova 2005)
- Reserve price important

Loss Averse Bidders

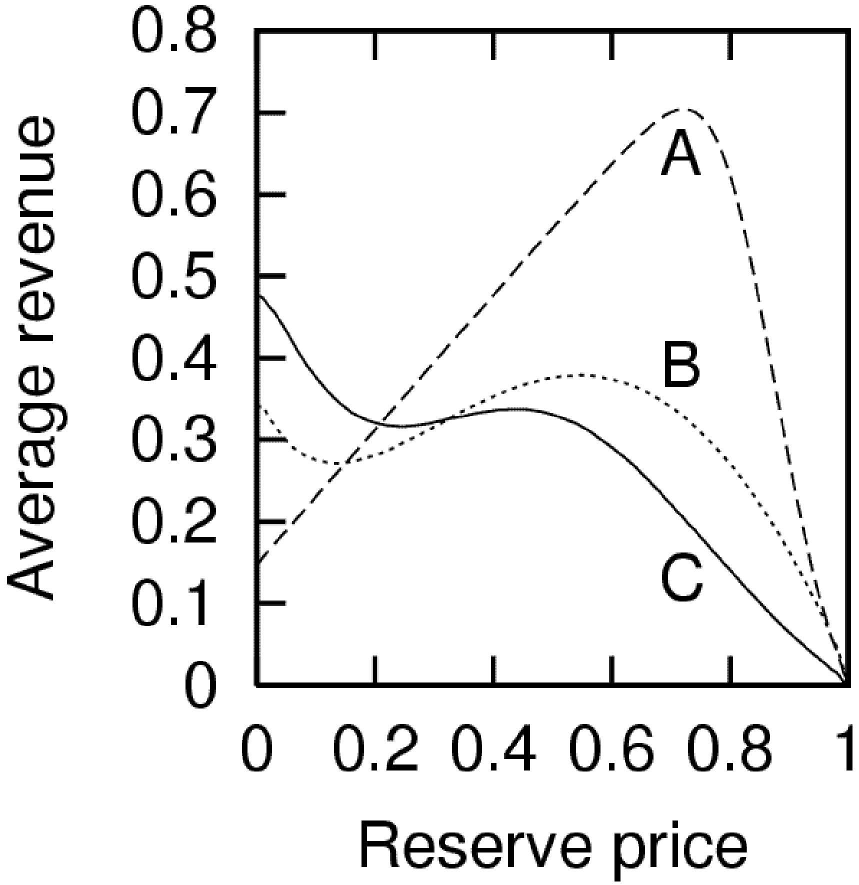

Auction Scenario

- Our seller has 1000 books to sell in auctions
  - series of English auctions with choice of reserve price
- The seller interacts with a population of bidders:
  - bidders characterized by valuation v, loss aversion
  - the population is characterized by distributions over v,
  - $0 < v < 1$; $1 <$ < 2.5
- Assume Gaussian distributions
  - mean of v chosen from [0, 1]; mean of   from [1, 2.5]
  - variances are 10^x, where x chosen from [-2, 1]
- 2 bidders per auction, following equilibrium

Individual populations:

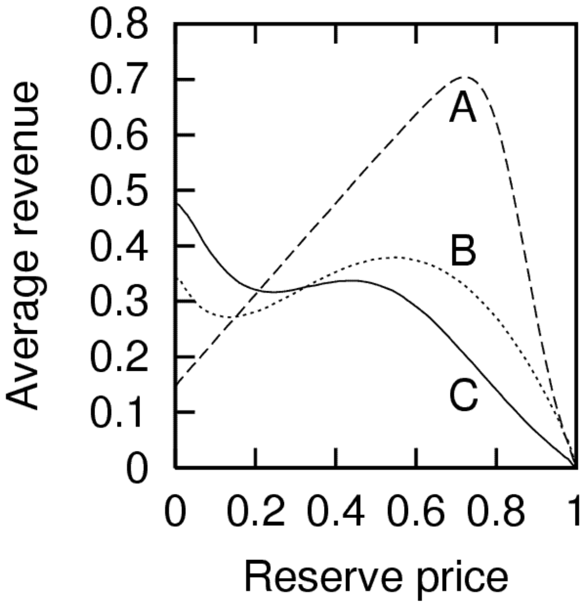

Average:

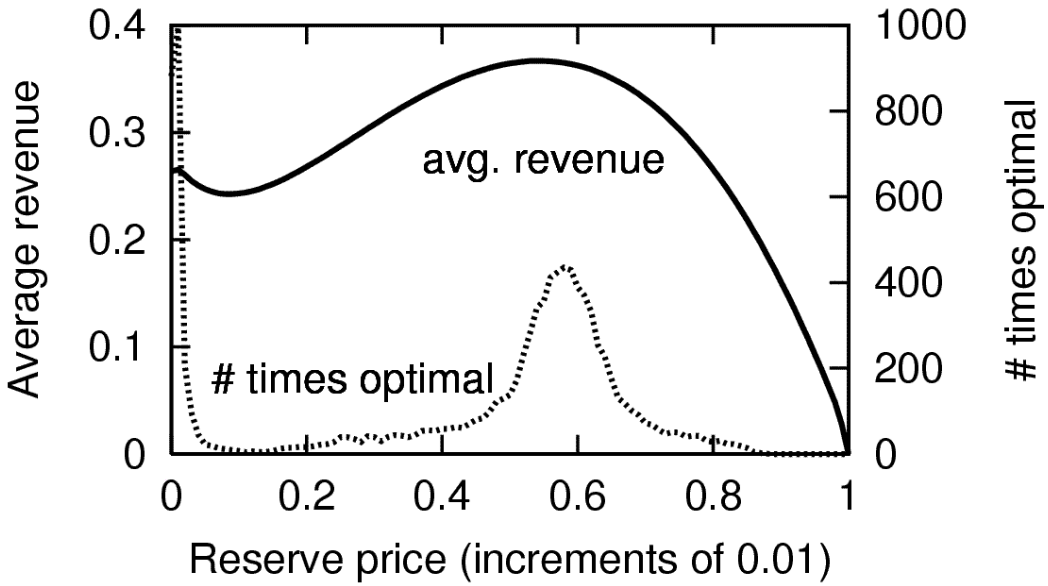

Overview

- Motivation
- Adaptive auction mechanisms
- Bidding scenario
- Adaptive mechanism implementation and results
- Incorporating predictions through metalearning
- Additional experiments

Adaptive Algorithm (Bandit)

- Discretize choices of reserve price (*k* choices)
- Results in a *k*-armed bandit problem
- Tradeoff between exploration and exploitation
- Sample averaging + softmax action selection:
  - Record avg_i and $count_i$ for each choice
  - Choose *i* with probability $(e^{avg_i/\tau})/(\sum_{j=1}^{k} e^{avg_j/\tau})$
  - controls exploration vs exploitation, often decreases

Adaptive Algorithm (Bandit)

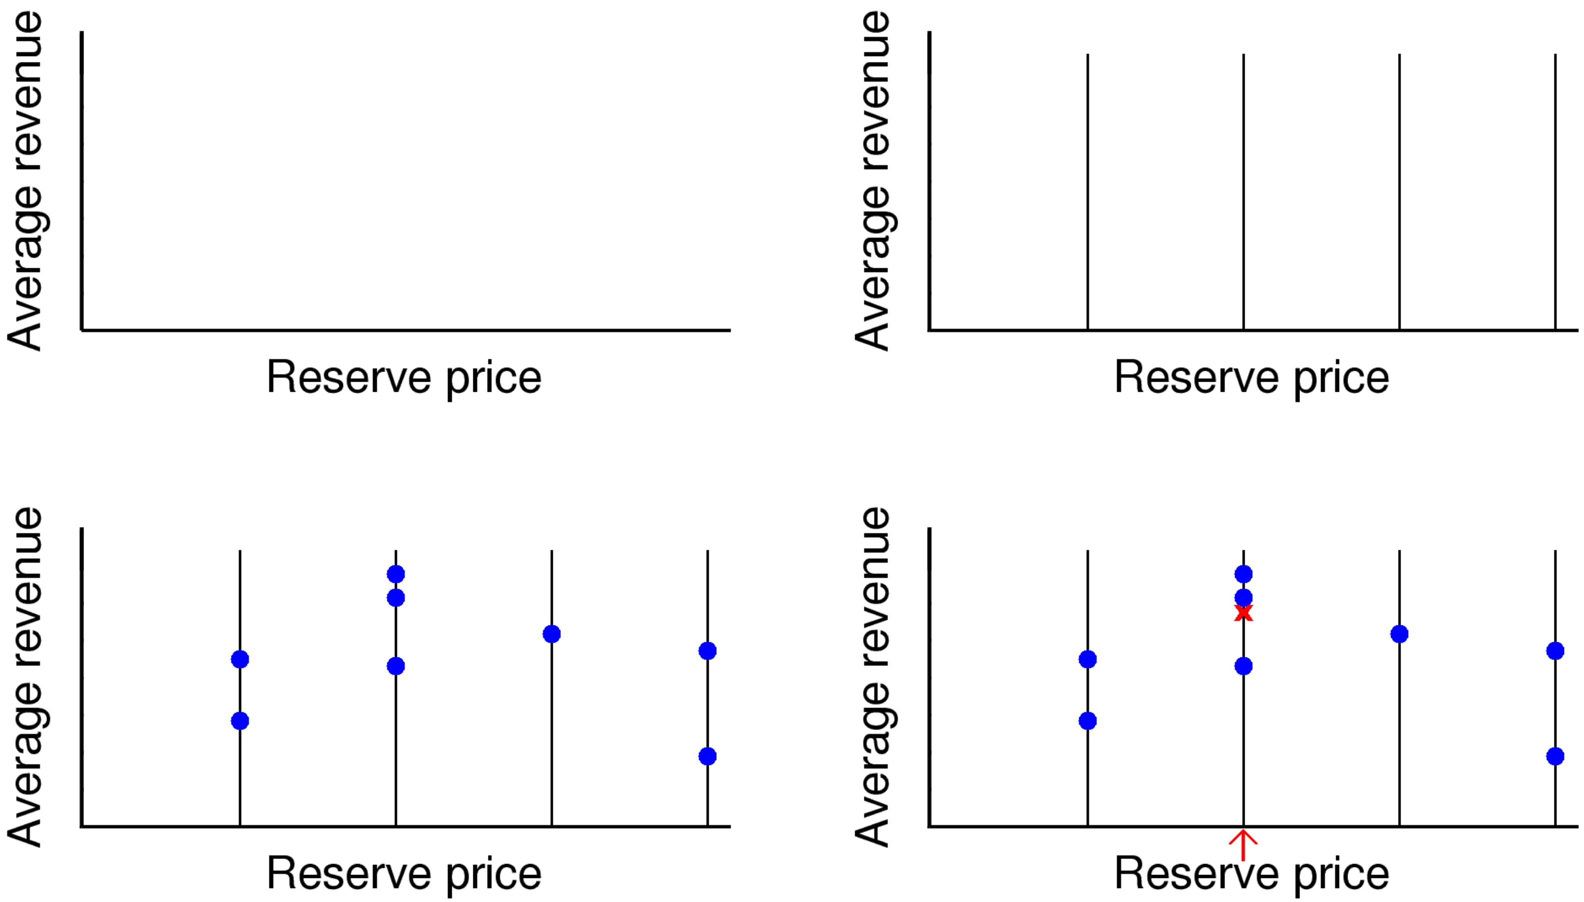

Adaptive Algorithm Parameters

- k (number of discrete choices)
- $_{start}$, $_{end}$ (decrease linearly over time)
- How to initialize values of avg_i and $count_i$?
  - optimistic initialization
- We choose these by hand:
  - k = 13
  - $_{start} = 0.1$, $_{end} = 0.01$
  - $avg_i = 0.6$, $count_i = 1$

Adaptive Algorithm (Regression)

- Bandit - restricts choices, assumes independence
- Solve by using regression:
  - Locally Weighted Quadratic Regression (instance based)
  - can estimate revenue at any point
  - considers all experience, uses a Gaussian weighting kernel
- Continue to discretize choices, but at high resolution
- Parameters nearly the same
  - need to choose kernel width (0.1)

Adaptive Algorithm (Regression)

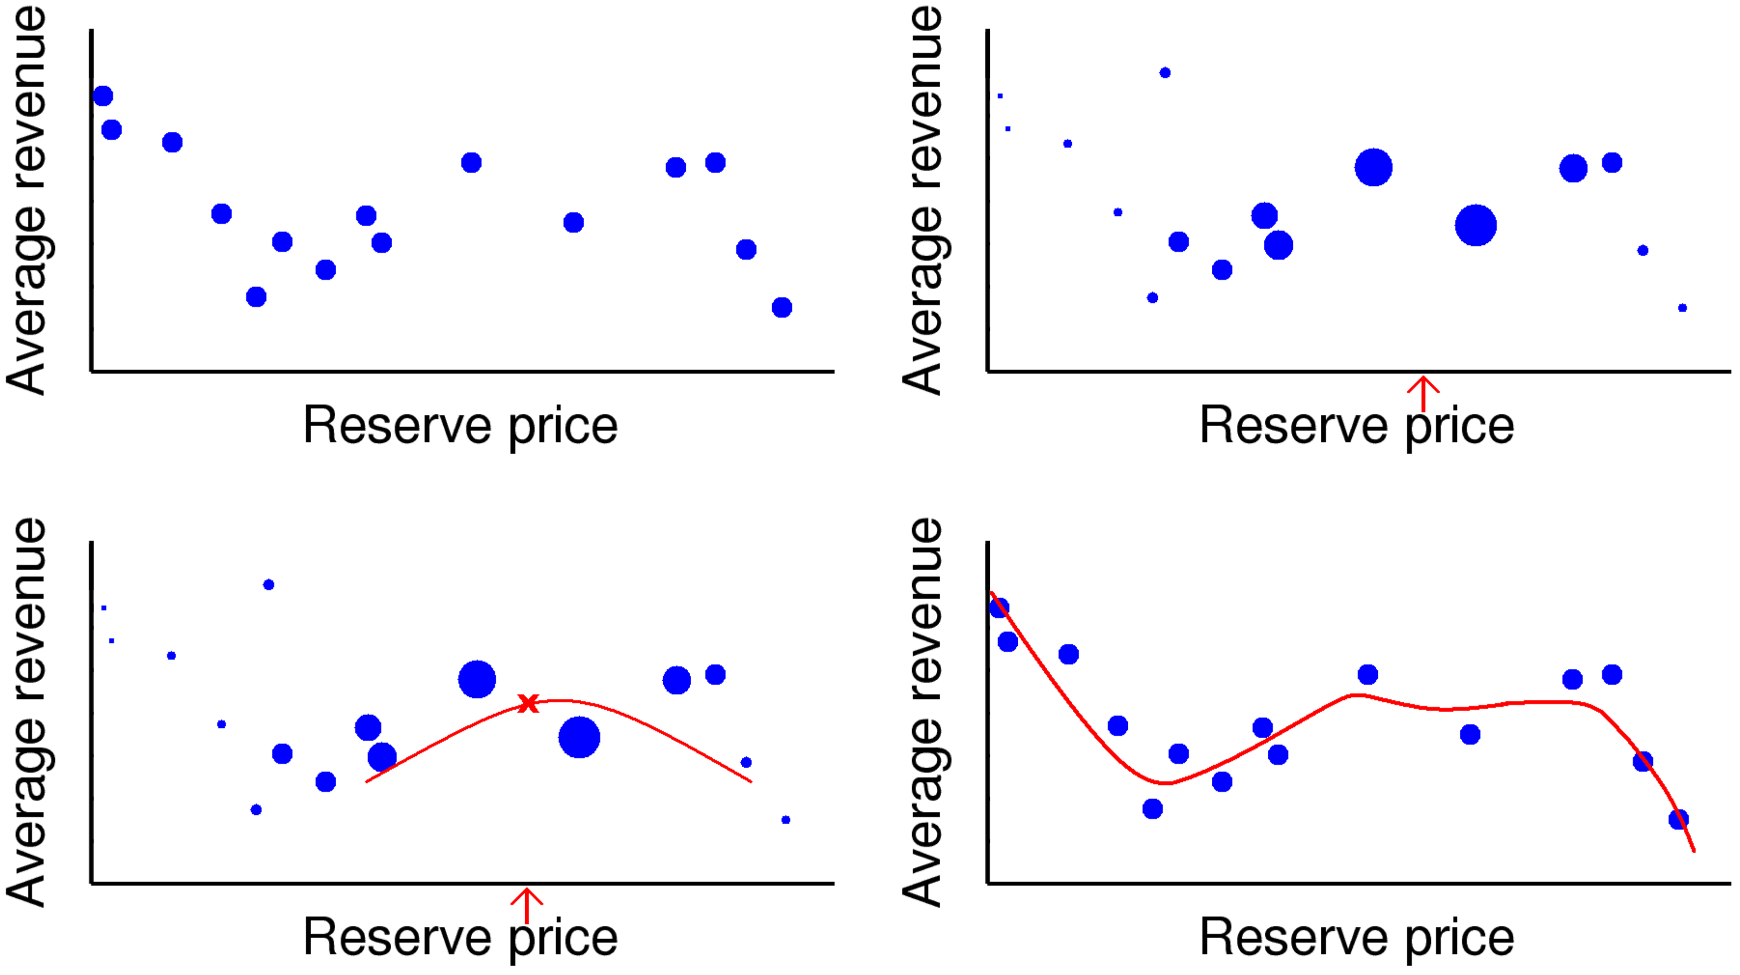

Results

- Average results over 10,000 generated bidder populations
- Significant with 99% confidence (paired t-tests)

Adaptive algorithm	Avg. revenue
best fixed reserve price (0.54)	0.367
bandit	0.374
regression	0.385

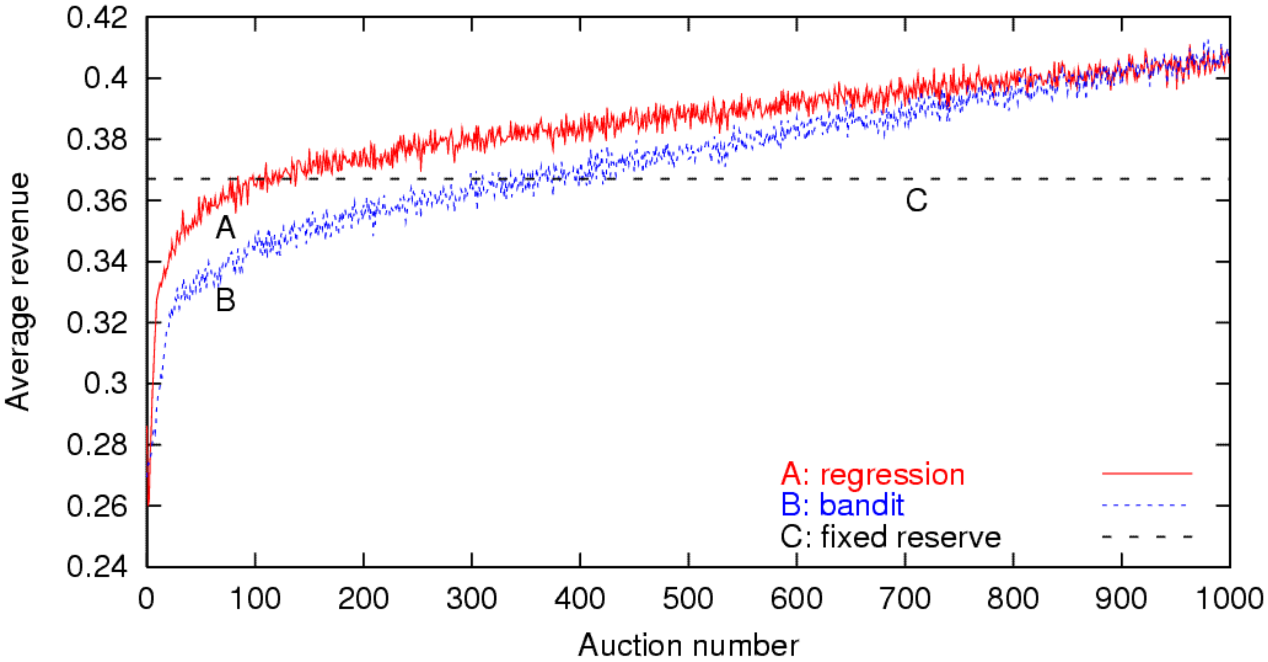

Overview

- Motivation
- Adaptive auction mechanisms
- Bidding scenario
- Adaptive mechanism implementation and results
- Incorporating predictions through metalearning
- Additional experiments

Taking Advantage of Predictions

- Adaptive mechanism requires no assumptions
- But what if reasonable predictions are possible?
- Example: selling a brand new book
  - could make guesses about bidder valuations, strategies
  - could consider books with similar author or subject

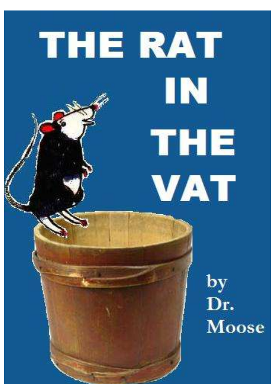

≈

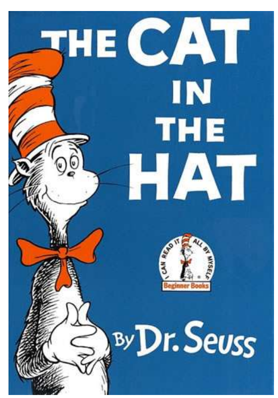

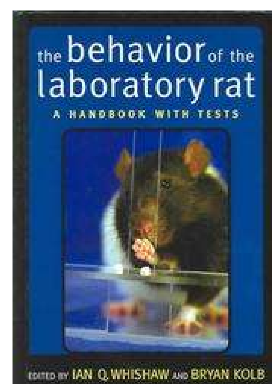

Taking Advantage of Predictions

- Seller can predict plausible bidder populations
- Adaptive mechanism should work well if correct

Metalearning

- Suppose seller can simulate bidder populations
- Choose an adaptive algorithm that is parameterized
- Search for optimal parameters in simulation
- An instance of metalearning

Metalearning

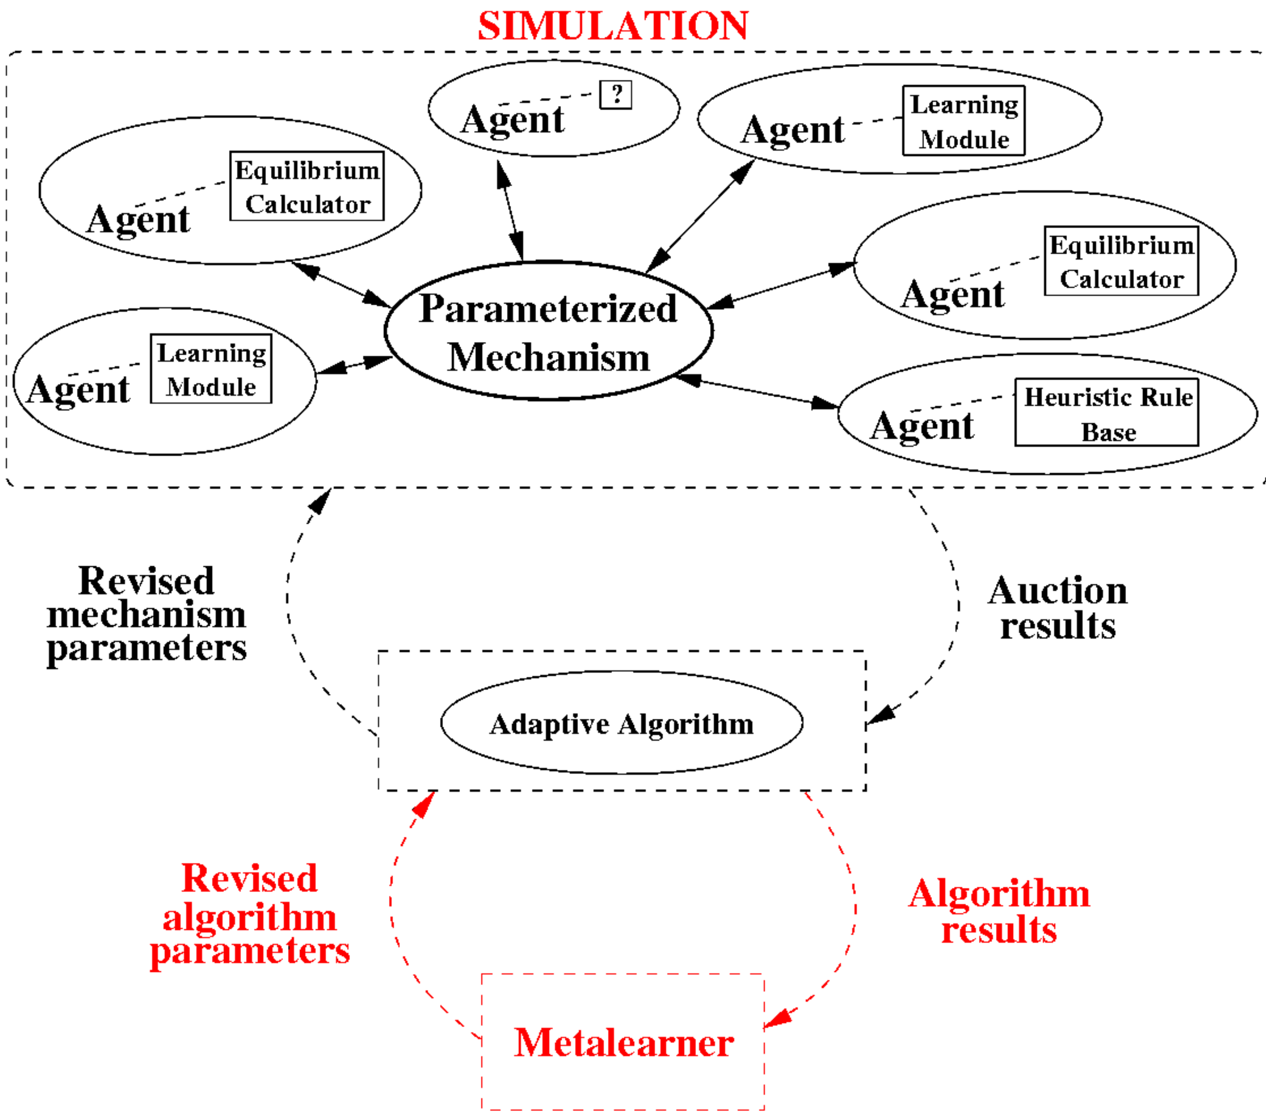

Simulation of Bidders by the Seller

- Suppose seller can predict possible populations (distributions of *v* and *)*
- Essentially a distribution over bidder populations
- Choose adaptive algorithm that performs best *with respect to this distribution*

Adaptive Parameters

- Now chosen through metalearning
- $_{start}$, $_{end}$
- Kernel width
- avg_i and $count_i$
  - optimistic initialization becomes initial experience

Parameter Search

- A stochastic optimization task
- Use Simultaneous Perturbation Stochastic Approximation (SPSA):
  - generate two estimates for slightly different parameters
  - move in direction of gradient
- Start with previously hand chosen parameters
- Time consuming, but done offline

Search Results

Bandit approach

$_{start}$ = .0423

$_{end}$ = .0077

Regression approach

$_{start}$ = .0081

$_{end}$ = .0013

kernel width = .138

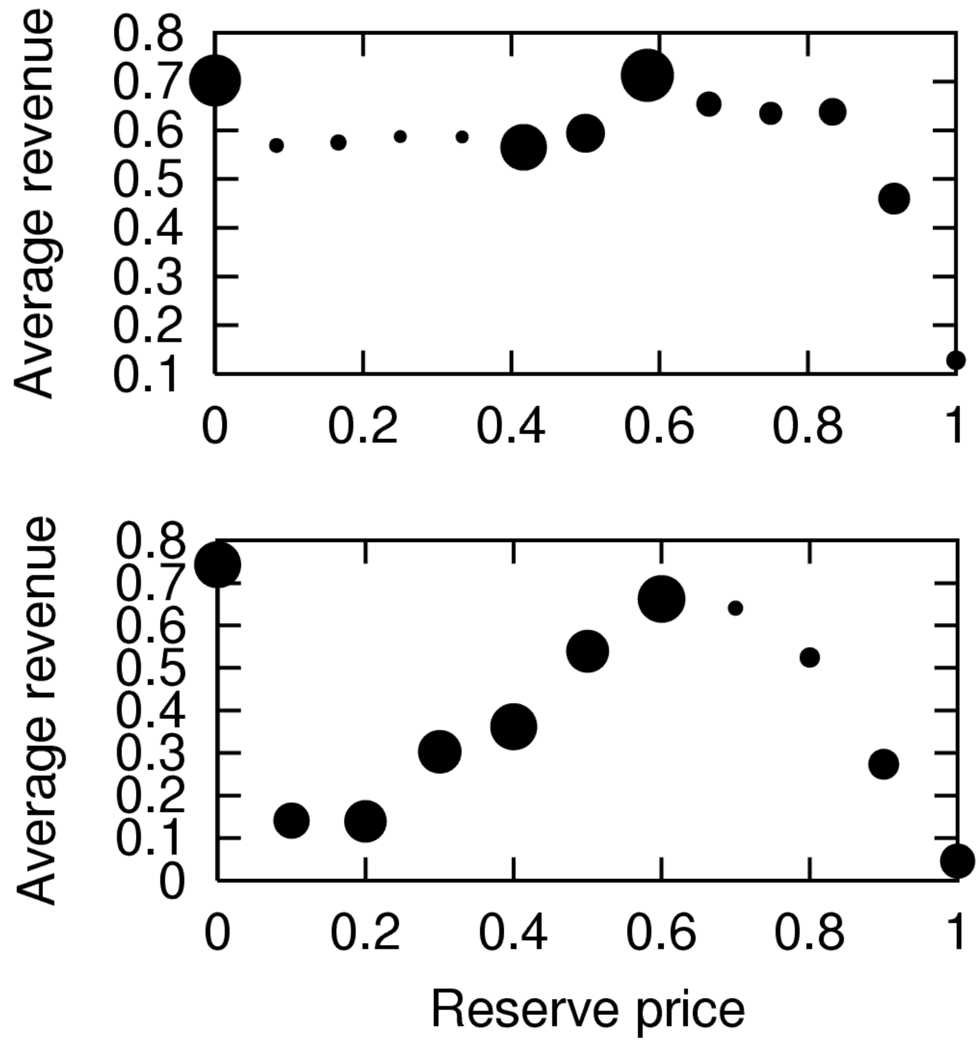

Results

- Average results over 10,000 populations drawn from predicted distribution
- Significant with 99% confidence (paired t-tests)

Adaptive algorithm	Avg. revenue
best fixed reserve price (0.54)	0.367
bandit, initial parameters	0.374
bandit, learned parameters	0.394
regression, initial parameters	0.385
regression, learned parameters	0.405

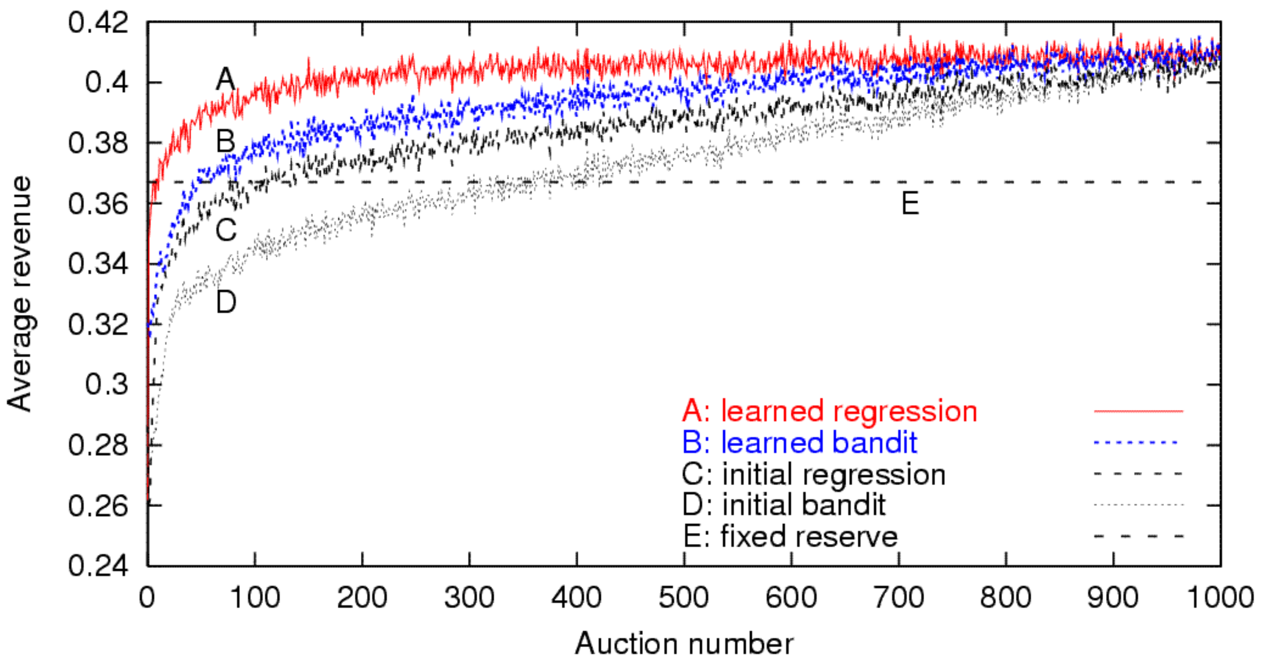

Overview

- Motivation
- Adaptive auction mechanisms
- Bidding scenario
- Adaptive mechanism implementation and results
- Incorporating predictions through metalearning
- Additional experiments

Questions

- Why not learn a model of the population?
- What if the population behaves unexpectedly? (different from simulated)
- What if the population changes over time?

Modeling the Population

- Bayesian approach
  - maintain probability distribution over possible populations (distributions of *v* and *)*
  - update after each new observation (auction result)
  - softmax action selection using expected revenues

Adaptive method	Avg. revenue
Bayesian approach	0.407
regression, initial parameters	0.385
regression, learned parameters	0.405

Unexpected Behavior

- Generate populations differently
  - before: mean of *v* in [0, 1]; mean of in [1, 2.5]
  - now: mean of *v* in [.3, .7]; mean of in [1.5, 2]

Adaptive method	Avg. revenue
Bayesian approach	0.414
regression, initial parameters	0.575
regression, learned parameters	0.593

Related Work

- Evolve ZIP traders and CDA together (Cliff 01)
- Evolve buyer and seller strategies and auction mechanism with genetic programming (Phelps et al. 02)
- Identify optimal price parameter of sealed bid auction for various bidder populations (Byde 03)

Future Work

- Encountered populations with unexpected behavior
- Non-stationary populations
- Learning populations
- Multiple mechanism parameters
- More sophisticated adaptive algorithms
- Evaluate on actual auction data

Conclusion

- Described design of adaptive auction mechanisms
- Experimented with a specific bidder scenario
- Adaptive mechanism outperforms fixed one
- Introduced metalearning approach
- Improve performance when predictions available