Automated Design of Robust Mechanisms

Michael Albert[1], Vincent Conitzer[1], Peter Stone[2]

[1]Duke University, [2]University of Texas at Austin

3rd Workshop on Algorithmic Game Theory and Data Science
June 26th, 2017

Previously appeared in AAAI17

Introduction - Revenue Efficient Mechanisms

- Standard mechanisms do very well with large numbers of bidders
  - VCG mechanism with $n+1$ bidders $\geq$ optimal revenue mechanism with n bidders, for IID bidders (Bulow and Klemperer 1996)
- For "thin" markets, must use knowledge of the distribution of bidders
  - Generalized second price auction with reserves (Myerson 1981)
- Thin markets are a large concern
  - Sponsored search with rare keywords or ad quality ratings
  - Of 19,688 reverse auctions by four governmental organizations in 2012, *one-third had only a single bidder* (GOA 2013)

Introduction - Revenue Efficient Mechanisms

- Standard mechanisms do very well with large numbers of bidders
  - VCG mechanism with $n+1$ bidders $\geq$ optimal revenue mechanism with n bidders, for IID bidders (Bulow and Klemperer 1996)
- For "thin" markets, must use knowledge of the distribution of bidders
  - Generalized second price auction with reserves (Myerson 1981)
- Thin markets are a large concern
  - Sponsored search with rare keywords or ad quality ratings
  - Of 19,688 reverse auctions by four governmental organizations in 2012, *one-third had only a single bidder* (GOA 2013)

Introduction - Revenue Efficient Mechanisms

- Standard mechanisms do very well with large numbers of bidders
  - VCG mechanism with $n+1$ bidders $\geq$ optimal revenue mechanism with n bidders, for IID bidders (Bulow and Klemperer 1996)
- For "thin" markets, must use knowledge of the distribution of bidders
  - Generalized second price auction with reserves (Myerson 1981)
- Thin markets are a large concern
  - Sponsored search with rare keywords or ad quality ratings
  - Of 19,688 reverse auctions by four governmental organizations in 2012, *one-third had only a single bidder* (GOA 2013)

Introduction - Correlated Distributions

- A common assumption in mechanism design is independent bidder valuations

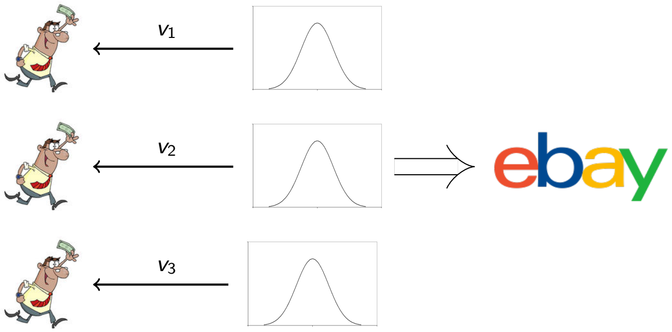

Introduction - Correlated Distributions

- This is not accurate for many settings
  - Oil drilling rights
  - Sponsored search auctions
  - Anything with resale value

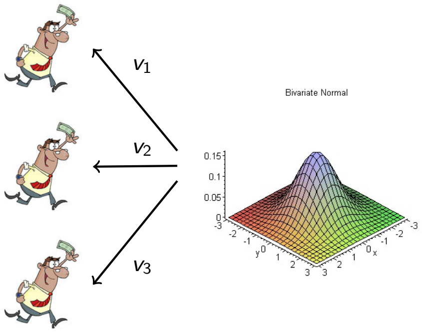

Introduction - Correlated Distributions

- Cremer and McLean (1985) demonstrates that full surplus extraction as revenue is possible for correlated valuation settings!

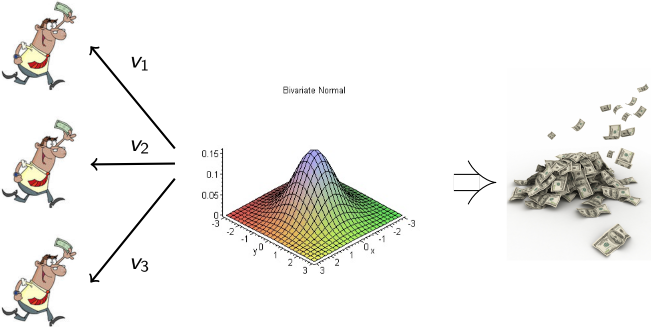

Contributions

How do we efficiently and robustly use distribution information?

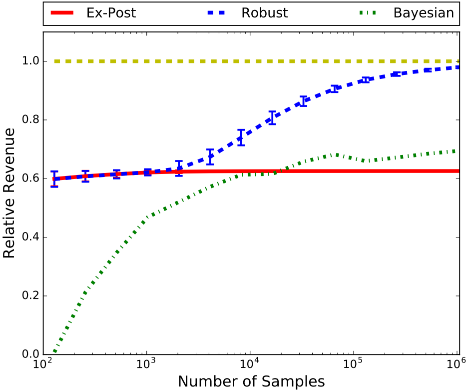

Problem Description

- A monopolistic seller with one item

- A single bidder with type $\theta \in \Theta$ and valuation $v(\theta)$
- An external signal $\omega \in \Omega$ and distribution $\pi(\theta, \omega)$

Problem Description

- A monopolistic seller with one item

- A single bidder with type $\theta \in \Theta$ and valuation $v(\theta)$

- An external signal $\omega \in \Omega$ and distribution $\pi(\theta, \omega)$

Problem Description

- A monopolistic seller with one item

- A single bidder with type $\theta \in \Theta$ and valuation $v(\theta)$

- An external signal $\omega \in \Omega$ and distribution $\pi(\theta, \omega)$

or

Definition: Ex-Post Individual Rationality (IR)

A mechanism $(\boldsymbol{p}, \boldsymbol{x})$ is *ex-post individually rational (IR)* if:

$$\forall \theta \in \Theta, \omega \in \Omega : U(\theta, \theta, \omega) \geq 0$$

Definition: Ex-Post Individual Rationality (IR)

A mechanism $(\boldsymbol{p}, \boldsymbol{x})$ is *ex-post individually rational (IR)* if:

$$\forall \theta \in \Theta, \omega \in \Omega : U(\theta, \theta, \omega) \geq 0$$

Definition: Bayesian Individual Rationality (IR)

A mechanism $(\boldsymbol{p}, \boldsymbol{x})$ is *Bayesian (or ex-interim) individually rational (IR)* if:

$$\forall \theta \in \Theta : \sum_{\omega \in \Omega} \pi(\omega|\theta) U(\theta, \theta, \omega) \geq 0$$

Definition: Ex-Post Individual Rationality (IR)

A mechanism $(\boldsymbol{p}, \boldsymbol{x})$ is *ex-post individually rational (IR)* if:

$$\forall \theta \in \Theta, \omega \in \Omega : U(\theta, \theta, \omega) \geq 0$$

Definition: Bayesian Individual Rationality (IR)

A mechanism $(\boldsymbol{p}, \boldsymbol{x})$ is *Bayesian (or ex-interim) individually rational (IR)* if:

$$\forall \theta \in \Theta : \sum_{\omega \in \Omega} \pi(\omega|\theta) U(\theta, \theta, \omega) \geq 0$$

Ex-Post IR Mechanisms $\subset$ *Bayesian IR Mechanisms*

Definition: Ex-Post Incentive Compatibility (IC)

A mechanism $(\boldsymbol{p}, \boldsymbol{x})$ is *ex-post incentive compatible (IC)* if:

$$\forall \theta, \theta' \in \Theta, \omega \in \Omega : U(\theta, \theta, \omega) \geq U(\theta, \theta', \omega)$$

Definition: Bayesian Incentive Compatibility (IC)

A mechanism $(\boldsymbol{p}, \boldsymbol{x})$ is *Bayesian incentive compatible (IC)* if:

$$\forall \theta, \theta' \in \Theta : \sum_{\omega \in \Omega} \pi(\omega|\theta) U(\theta, \theta, \omega) \geq \sum_{\omega \in \Omega} \pi(\omega|\theta) U(\theta, \theta', \omega)$$

Ex-Post IC Mechanisms $\subset$ *Bayesian IC Mechanisms*

Definition: Optimal Ex-Post Mechanisms

A mechanism $(\boldsymbol{p}, \boldsymbol{x})$ is an *optimal ex-post mechanism* if under the constraint of ex-post individual rationality and ex-post incentive compatibility it maximizes the following:

$$\sum_{\theta,\omega} \boldsymbol{x}(\theta,\omega)\boldsymbol{\pi}(\theta,\omega) \tag{1}$$

Definition: Optimal Bayesian Mechanism

A mechanism that maximizes (1) under the constraint of Bayesian individual rationality and Bayesian incentive compatibility is an *optimal Bayesian mechanism*.

Ex-Post Revenue $\leq$ *Bayesian Revenue*

Review of Bayesian Mechanism Design

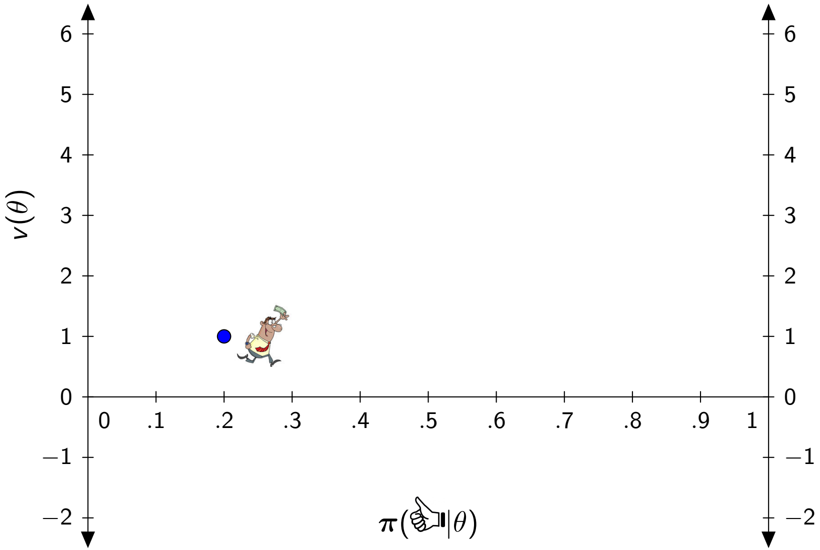

Review of Bayesian Mechanism Design

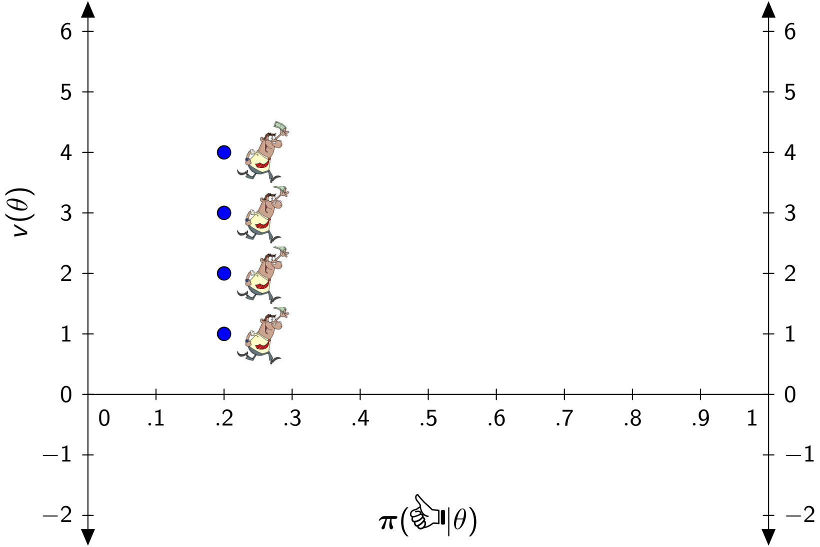

Review of Bayesian Mechanism Design

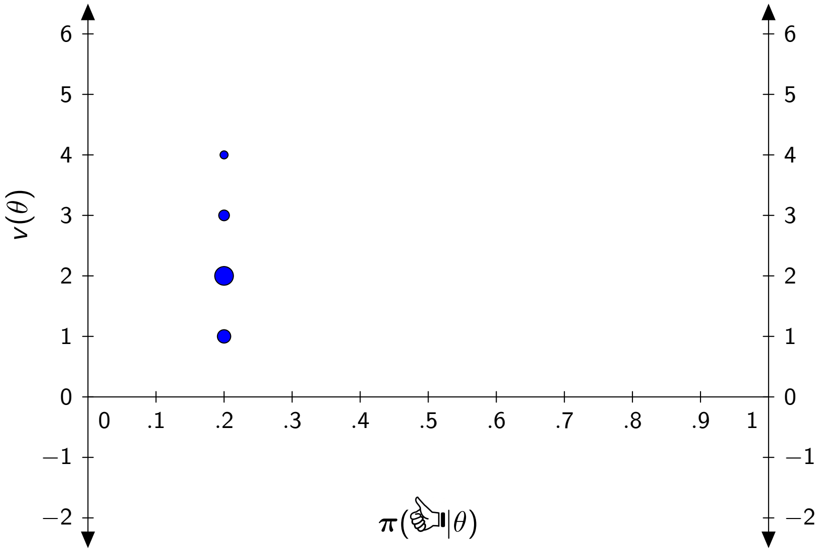

Review of Bayesian Mechanism Design

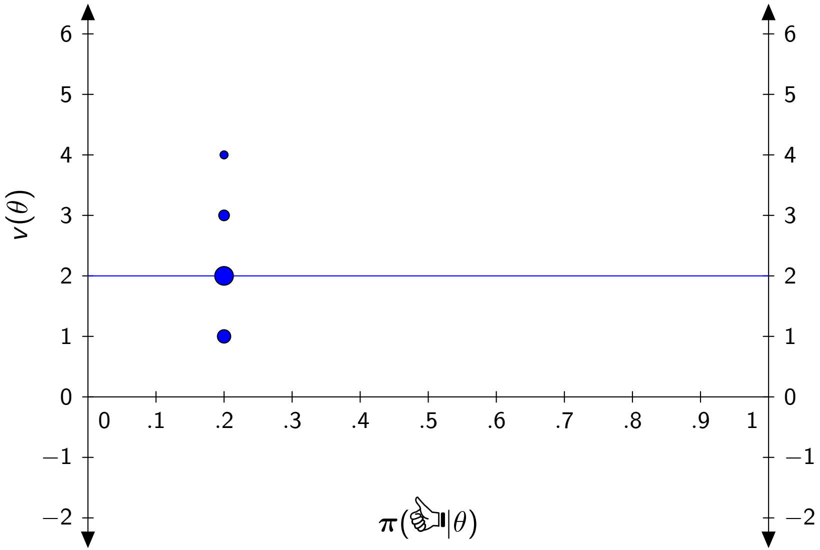

Review of Bayesian Mechanism Design

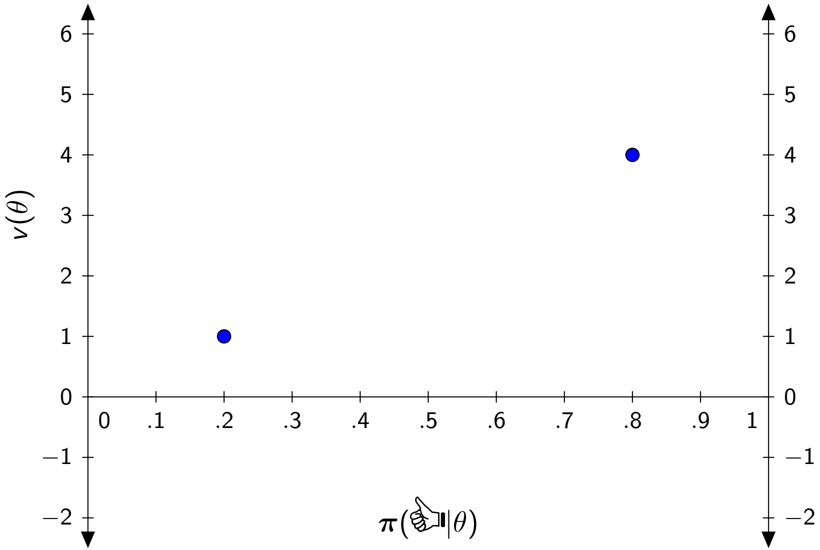

Review of Bayesian Mechanism Design

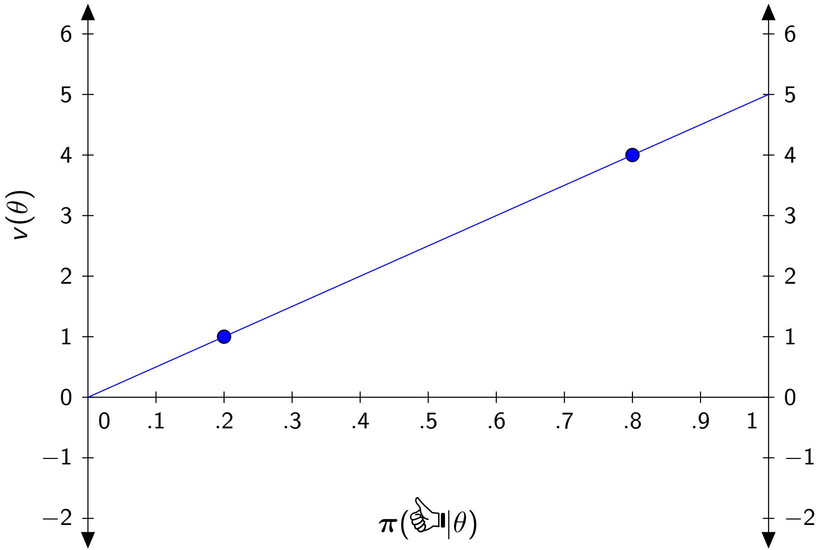

Distribution Uncertainty

What if the distribution isn't well known?

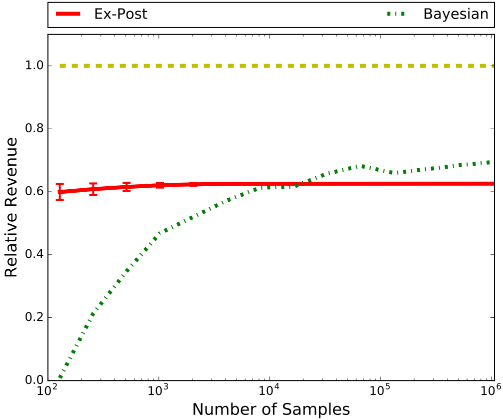

Robust Mechanism Design

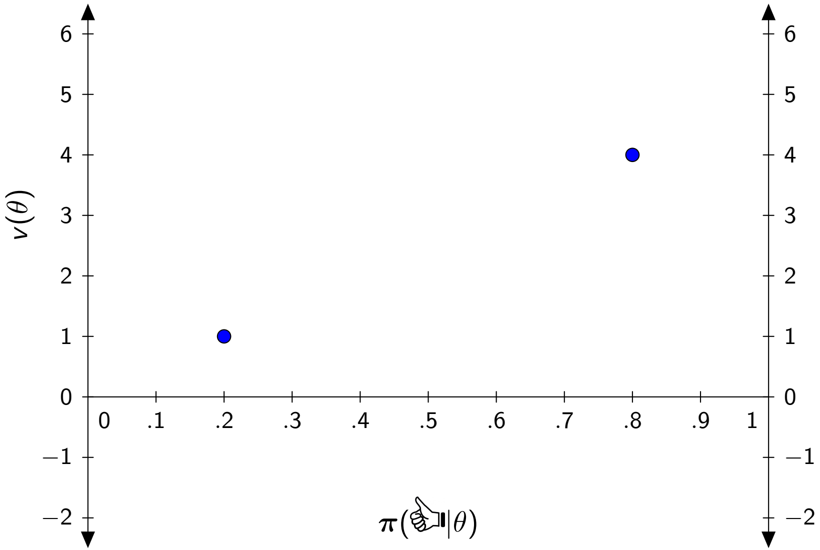

Robust Mechanism Design

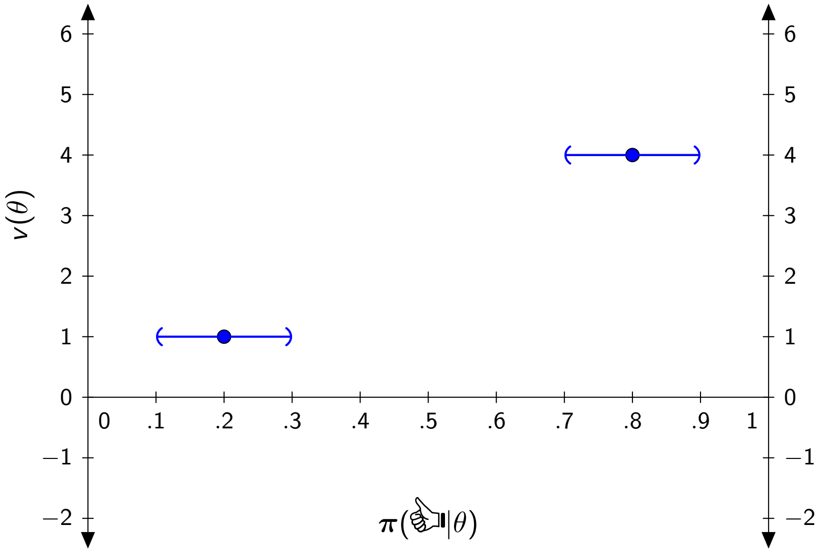

Robust Mechanism Design

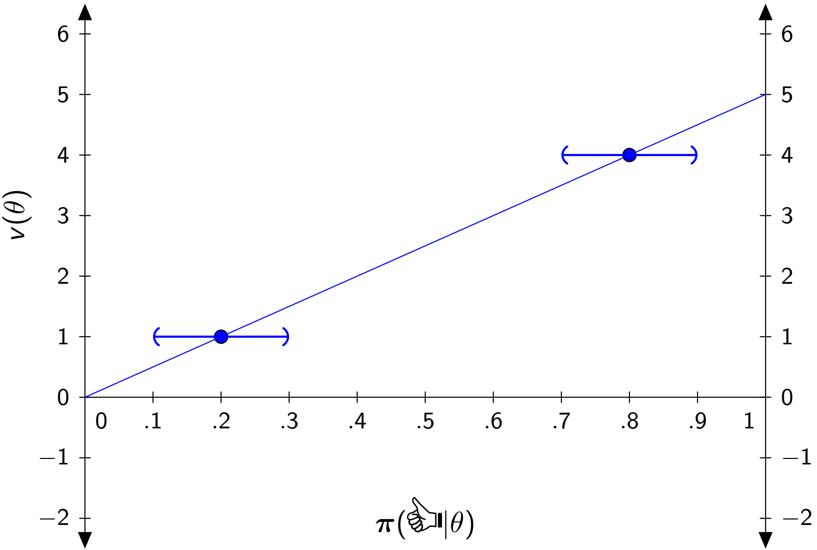

Robust Mechanism Design

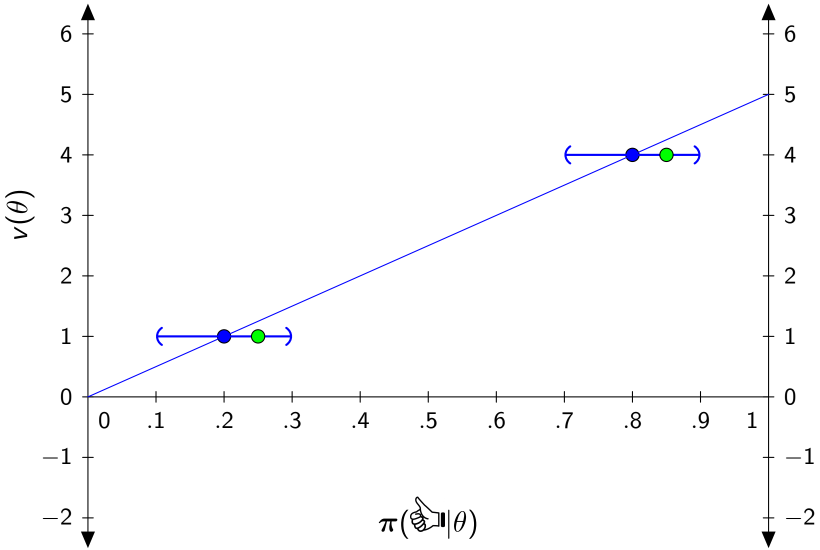

Robust Mechanism Design

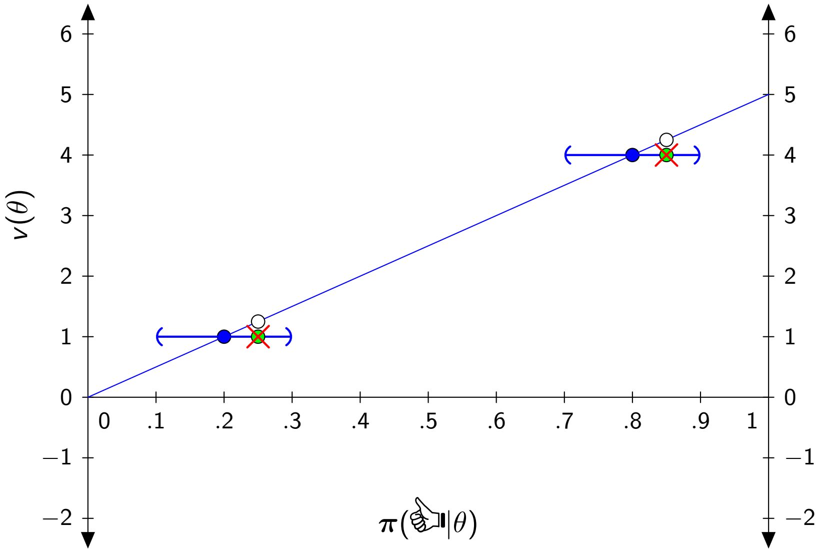

Robust Mechanism Design

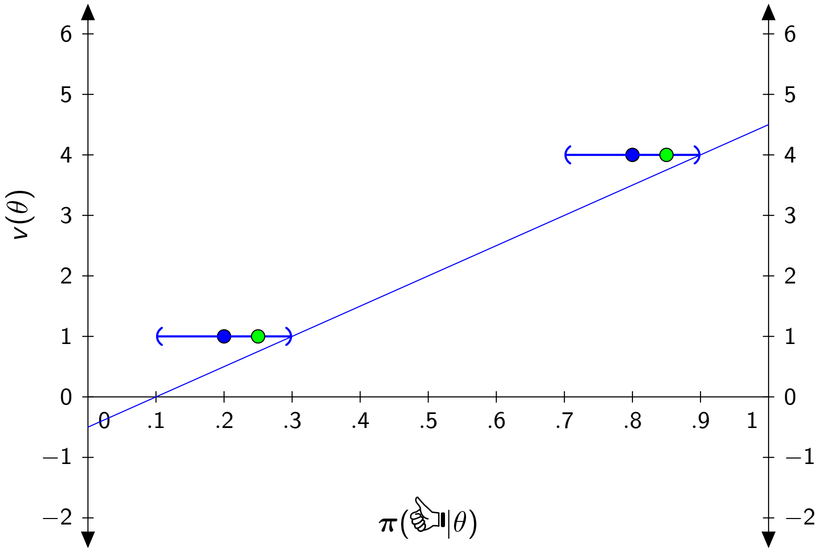

Robust Mechanism Design

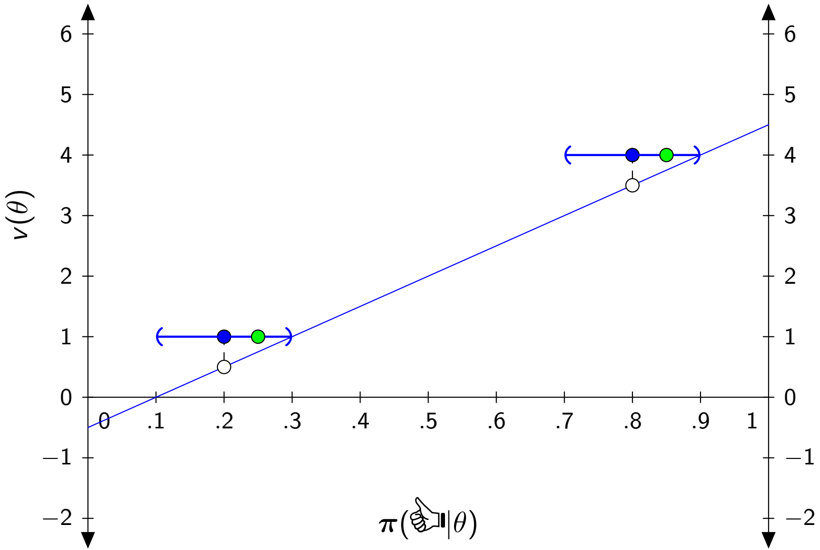

Consistent Distributions

Definition: Set of Consistent Distributions

Let $P(A)$ be the set of probability distributions over A. Then the space of all probability distributions over $\Theta \times \Omega$ can be represented as $P(\Theta \times \Omega)$. A subset $\mathcal{P}(\hat{\pi}) \subseteq P(\Theta \times \Omega)$ is a *consistent set of distributions* for the estimated distribution $\hat{\pi}$ if the true distribution, π, is guaranteed to be in $\mathcal{P}(\hat{\pi})$ and $\hat{\pi} \in \mathcal{P}(\hat{\pi})$.

Robust IR and IC

Definition: Robust Individual Rationality

A mechanism is *robust individually rational* for estimated bidder distribution $\hat{\boldsymbol{\pi}}$ and consistent set of distributions $\mathcal{P}(\hat{\boldsymbol{\pi}})$ if for all $\theta \in \Theta$ and $\boldsymbol{\pi} \in \mathcal{P}(\hat{\boldsymbol{\pi}})$,

$$\sum_{\omega \in \Omega} \boldsymbol{\pi}(\omega|\theta) U(\theta, \boldsymbol{\pi}, \theta, \boldsymbol{\pi}, \omega) \geq 0$$

Robust IR and IC

Definition: Robust Individual Rationality

A mechanism is *robust individually rational* for estimated bidder distribution $\hat{\boldsymbol{\pi}}$ and consistent set of distributions $\mathcal{P}(\hat{\boldsymbol{\pi}})$ if for all $\theta \in \Theta$ and $\boldsymbol{\pi} \in \mathcal{P}(\hat{\boldsymbol{\pi}})$,

$$\sum_{\omega \in \Omega} \boldsymbol{\pi}(\omega|\theta) U(\theta, \boldsymbol{\pi}, \theta, \boldsymbol{\pi}, \omega) \geq 0$$

Definition: Robust Incentive Compatibility

A mechanism is *robust incentive compatible* for estimated bidder distribution $\hat{\boldsymbol{\pi}}$ and consistent set of distributions $\mathcal{P}(\hat{\boldsymbol{\pi}})$ if for all $\theta, \theta' \in \Theta$ and $\boldsymbol{\pi}, \boldsymbol{\pi}' \in \mathcal{P}(\hat{\boldsymbol{\pi}})$,

$$\sum_{\omega \in \Omega} \boldsymbol{\pi}(\omega|\theta) U(\theta, \boldsymbol{\pi}, \theta, \boldsymbol{\pi}, \omega) \geq \sum_{\omega \in \Omega} \boldsymbol{\pi}(\omega|\theta) U(\theta, \boldsymbol{\pi}, \theta', \boldsymbol{\pi}', \omega)$$

Robust IR and IC

Heirarchy of Individual Rationality

$$\textit{Ex-Post IR} \subseteq \textit{Robust IR} \subseteq \textit{Bayesian IR}$$

Heirarchy of Incentive Compatibility

$$\textit{Ex-Post IC} \subseteq \textit{Robust IC} \subseteq \textit{Bayesian IC}$$

Definition: Optimal Restricted Robust Mechanism

The *optimal restricted robust mechanism* given an estimated distribution $\hat{\pi}$ and a consistent set of distributions $\mathcal{P}(\hat{\pi})$ is a mechanism dependent only on the reported type and exernal signal that maximizes the following objective:

$$\sum_{\theta,\omega} \hat{\pi}(\theta,\omega)\boldsymbol{x}(\theta,\omega)$$

while satisfying robust IC and IR with respect to $\mathcal{P}(\hat{\pi})$.

Definition: Optimal Restricted Robust Mechanism

The *optimal restricted robust mechanism* given an estimated distribution $\hat{\pi}$ and a consistent set of distributions $\mathcal{P}(\hat{\pi})$ is a mechanism dependent only on the reported type and exernal signal that maximizes the following objective:

$$\sum_{\theta,\omega} \hat{\pi}(\theta,\omega)\boldsymbol{x}(\theta,\omega)$$

while satisfying robust IC and IR with respect to $\mathcal{P}(\hat{\pi})$.

Heirarchy of Revenue

Ex-Post Mechanism $\leq$ *Robust Mechanism* $\leq$ *Bayesian Mechanism*

Polynomial Time Algorithm

Assumption: Polyhedral Consistent Set

The set $\mathcal{P}(\hat{\pi})$ can be characterized as an $n-$polyhedron, where n is polynomial in the number of bidder types and external signals.

Polynomial Time Algorithm

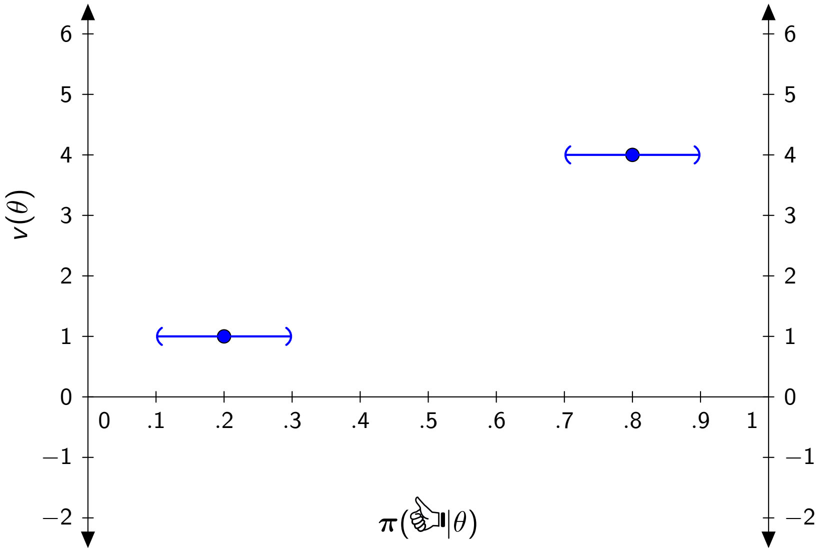

Polynomial Time Algorithm

Assumption: Polyhedral Consistent Set

The set $\mathcal{P}(\hat{\pi})$ can be characterized as an $n-$polyhedron, where n is polynomial in the number of bidder types and external signals.

Theorem: Polynomial Complexity of the Optimal Restricted Robust Mechanism

If $\mathcal{P}(\hat{\pi})$ satisfies the above assumption, the optimal restricted robust mechanism can be calculated in time polynomial in the number of types of the bidder and external signal.

Varying Between Ex-Post and Bayesian

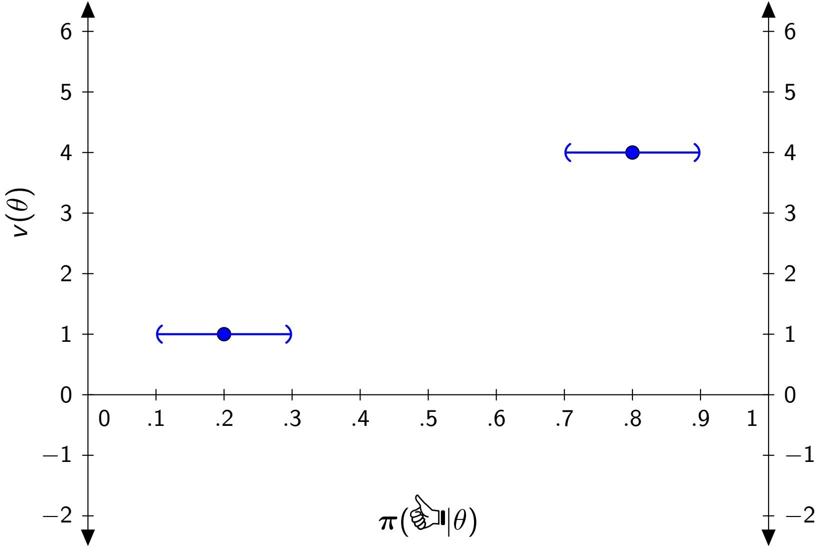

Varying Between Ex-Post and Bayesian

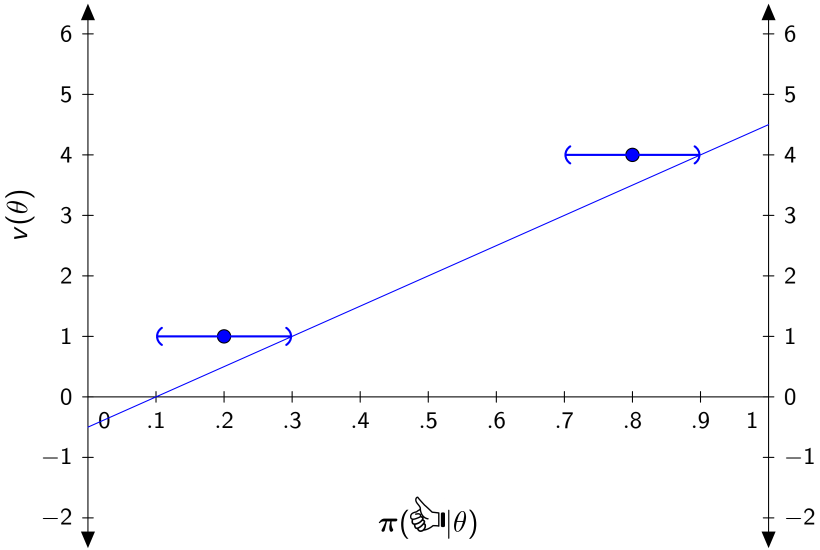

Varying Between Ex-Post and Bayesian

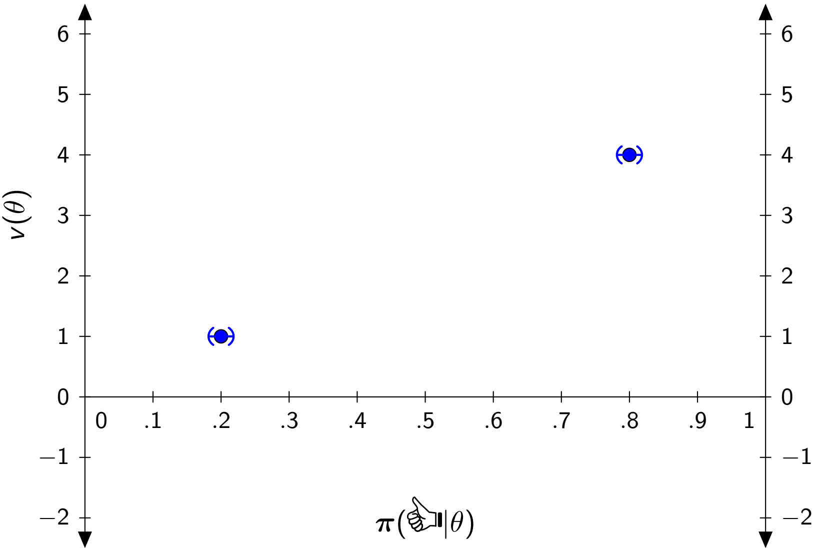

Varying Between Ex-Post and Bayesian

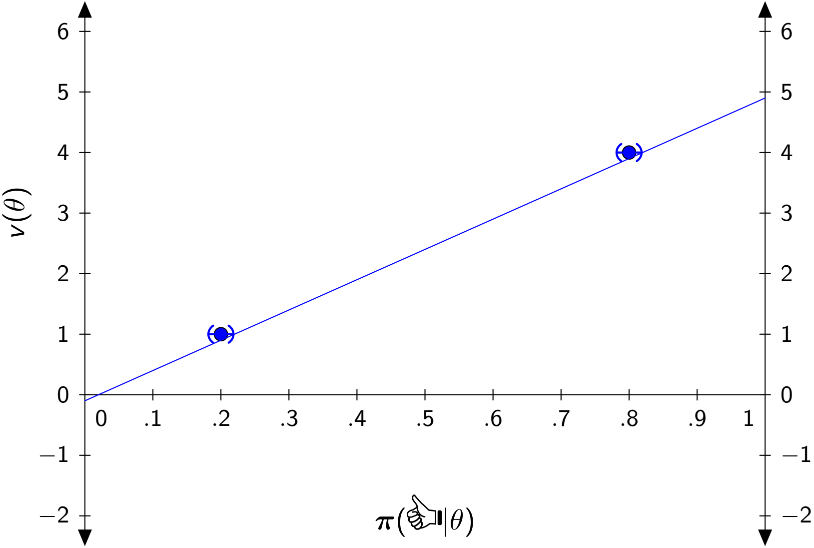

Varying Between Ex-Post and Bayesian

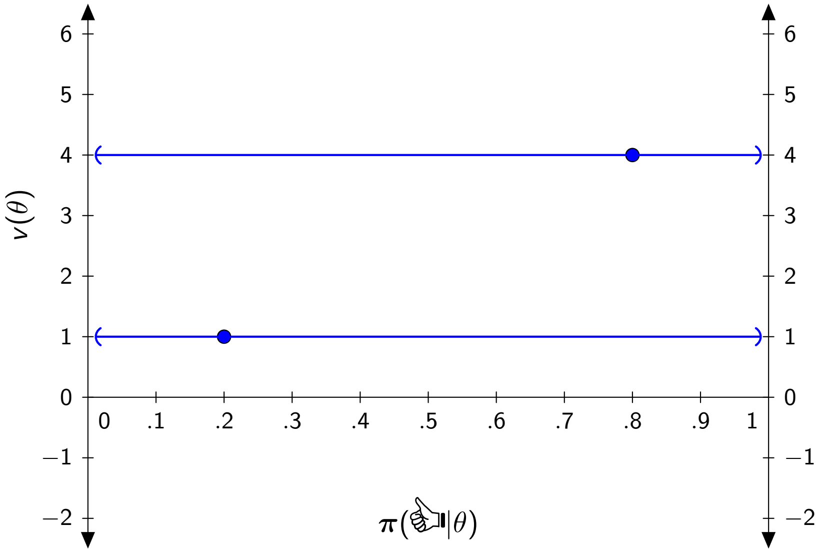

Varying Between Ex-Post and Bayesian

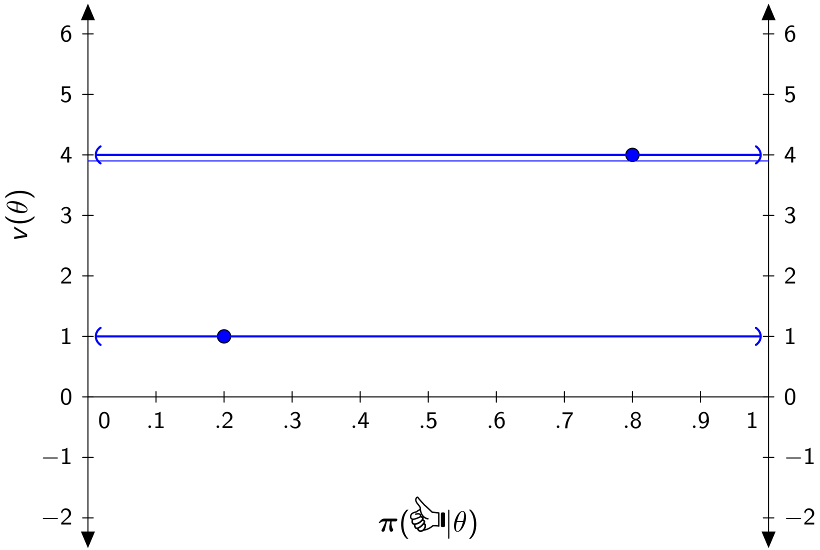

ϵ-Robust Mechanism Design

Robust is not sufficient

- All results and intuition for restricted robust mechanism design carries over to restricted ϵ-robust mechanism design

ϵ-Robust Mechanism Design

Robust is not sufficient

Definition: Set of ϵ-Consistent Distributions

A subset $\mathcal{P}_\epsilon(\hat{\pi}) \subseteq P(\Theta \times \Omega)$ is an *ϵ-consistent set of distributions* for the estimated distribution $\hat{\pi}$ if the true distribution, π, is in $\mathcal{P}_\epsilon(\hat{\pi})$ with probability $1 - \epsilon$ and $\hat{\pi} \in \mathcal{P}_\epsilon(\hat{\pi})$.

- All results and intuition for restricted robust mechanism design carries over to restricted ϵ-robust mechanism design

ϵ-Robust Mechanism Design

Robust is not sufficient

Definition: Set of ϵ-Consistent Distributions

A subset $\mathcal{P}_\epsilon(\hat{\pi}) \subseteq P(\Theta \times \Omega)$ is an *ϵ-consistent set of distributions* for the estimated distribution $\hat{\pi}$ if the true distribution, π, is in $\mathcal{P}_\epsilon(\hat{\pi})$ with probability $1 - \epsilon$ and $\hat{\pi} \in \mathcal{P}_\epsilon(\hat{\pi})$.

- All results and intuition for restricted robust mechanism design carries over to restricted ϵ-robust mechanism design

Experiments

- True distribution is discretized bivariate normal distribution
- Sample from the true distribution N times
- Use Bayesian methods to estimate the distribution
- Calculate empirical confidence intervals for elements of the distribution
- Parameters unless otherwise specified:
  - Correlation $= .5$
  - $\epsilon = .05$
  - $\Theta = \{1, 2, ..., 10\}$
  - $|\Omega| = 10$
  - $v(\theta) = \theta$

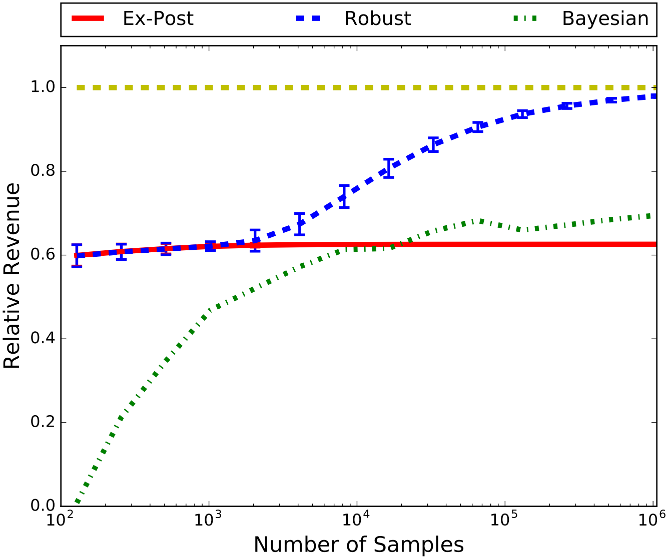

Ex-Post
Robust
Bayesian
1.0
0.8
0.6
0.4
0.2
0.0
10^2
10^3
10^4
10^5
10^6
Relative Revenue
Number of Samples

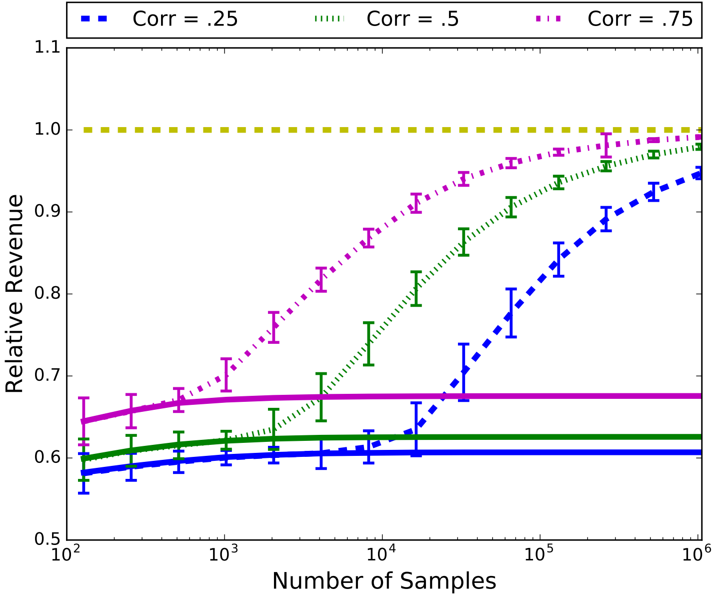
Corr = .25
Corr = .5
Corr = .75
1.1
1.0
0.9
0.8
0.7
0.6
0.5
Relative Revenue
10^2
10^3
10^4
10^5
10^6
Number of Samples

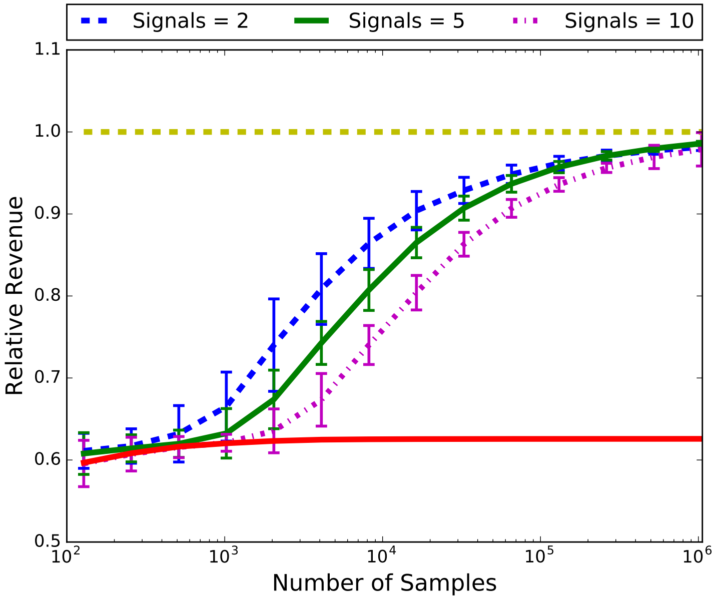
Signals = 2
Signals = 5
Signals = 10
Relative Revenue
1.1
1.0
0.9
0.8
0.7
0.6
0.5
10^2
10^3
10^4
10^5
10^6
Number of Samples

Related Work

- Uncertainty in Mechanism Design (Lopomo, Rigotti, and Shannon 2009, 2011)
- Automated Mechanism Design (Conitzer and Sandholm 2002, 2004; Guo and Conitzer 2010; Sandholm and Likhodedov 2015)
- Robust Optimization (Bertsimas and Sim 2004; Aghassi and Bertsimas 2006)
- Learning Bidder Distribution (Elkind 2007, Fu et al 2014, Blume et. al. 2015, Morgenstern and Roughgarden 2015)
- Simple vs. Optimal Mechanisms (Bulow and Klemperer 1996; Hartline and Roughgarden 2009)

Thank you for listening to my presentation.
Questions?