Mechanism Design with Unknown Correlated Distributions: Can We Learn Optimal Mechanisms?

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

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

May 10th, 2017

Introduction

- Auctions are one of the fundamental tools of the modern economy
  - In 2012, four government agencies purchased **$800 million** through reverse auctions (Government Accountability Office 2013)
  - In 2014, NASA awarded contracts to Boeing and Space-X worth **$4.2 billion** and **$2.6 billion** through an auction process (NASA 2014)
  - In 2016, **$72.5 billion** of ad revenue generated through auctions (IAB 2017)
  - The FCC spectrum auction just allocated **$20 billion** worth of broadcast spectrum

Introduction

- Auctions are one of the fundamental tools of the modern economy
    - In 2012, four government agencies purchased **$800 million** through reverse auctions (Government Accountability Office 2013)
    - In 2014, NASA awarded contracts to Boeing and Space-X worth **$4.2 billion** and **$2.6 billion** through an auction process (NASA 2014)
    - In 2016, **$72.5 billion** of ad revenue generated through auctions (IAB 2017)
    - The FCC spectrum auction just allocated **$20 billion** worth of broadcast spectrum

Introduction

- Auctions are one of the fundamental tools of the modern economy
  - In 2012, four government agencies purchased **$800 million** through reverse auctions (Government Accountability Office 2013)
  - In 2014, NASA awarded contracts to Boeing and Space-X worth **$4.2 billion** and **$2.6 billion** through an auction process (NASA 2014)
  - In 2016, **$72.5 billion** of ad revenue generated through auctions (IAB 2017)
  - The FCC spectrum auction just allocated **$20 billion** worth of broadcast spectrum

Introduction

- Auctions are one of the fundamental tools of the modern economy
  - In 2012, four government agencies purchased **$800 million** through reverse auctions (Government Accountability Office 2013)
  - In 2014, NASA awarded contracts to Boeing and Space-X worth **$4.2 billion** and **$2.6 billion** through an auction process (NASA 2014)
  - In 2016, **$72.5 billion** of ad revenue generated through auctions (IAB 2017)
  - The FCC spectrum auction just allocated **$20 billion** worth of broadcast spectrum

Introduction

- Auctions are one of the fundamental tools of the modern economy
  - In 2012, four government agencies purchased **$800 million** through reverse auctions (Government Accountability Office 2013)
  - In 2014, NASA awarded contracts to Boeing and Space-X worth **$4.2 billion** and **$2.6 billion** through an auction process (NASA 2014)
  - In 2016, **$72.5 billion** of ad revenue generated through auctions (IAB 2017)
  - The FCC spectrum auction just allocated **$20 billion** worth of broadcast spectrum

Introduction

- Auctions are one of the fundamental tools of the modern economy
  - In 2012, four government agencies purchased **$800 million** through reverse auctions (Government Accountability Office 2013)
  - In 2014, NASA awarded contracts to Boeing and Space-X worth **$4.2 billion** and **$2.6 billion** through an auction process (NASA 2014)
  - In 2016, **$72.5 billion** of ad revenue generated through auctions (IAB 2017)
  - The FCC spectrum auction just allocated **$20 billion** worth of broadcast spectrum

It is important that the mechanisms we use are revenue optimal!

Introduction

- 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

- 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

- 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

- A common assumption in mechanism design is independent bidder valuations

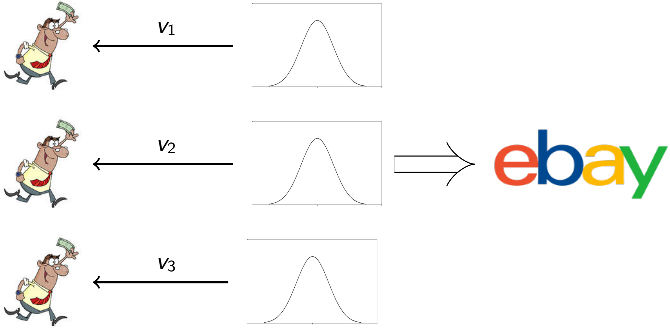

Introduction

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

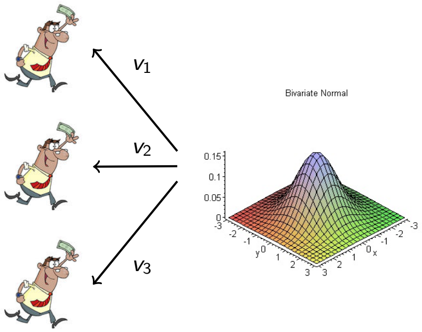

Introduction

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

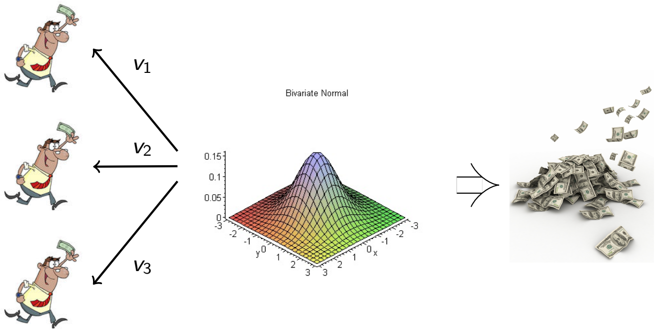

Introduction

- What if we don't know the distribution though?

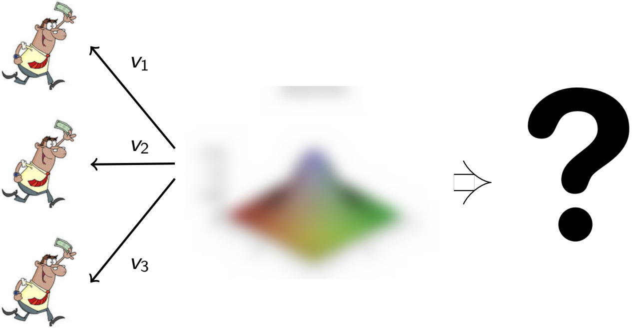

Introduction

- Fu et. al. 2014 indicate that it is still easy if we have a finite set of potential distributions!

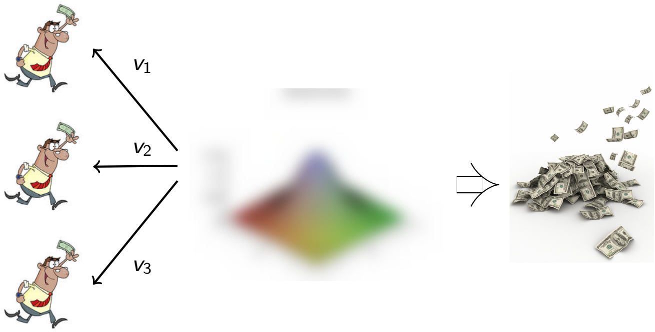

Introduction

- What if we have an infinite set of distributions?

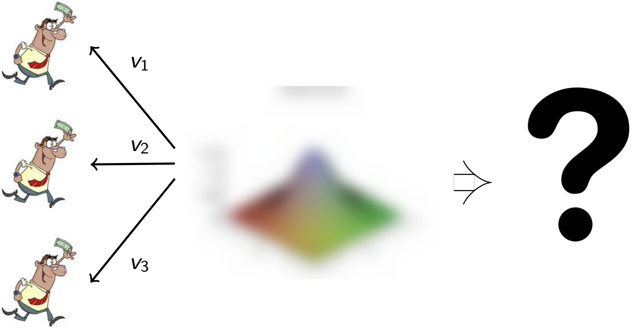

Contribution

In order to effectively implement mechanisms that take advantage of correlation, there needs to be a lot of correlation.

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 $\pi(\theta, \omega) \in \Delta(\Theta \times \Omega)$

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 $\pi(\theta, \omega) \in \Delta(\Theta \times \Omega)$

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 $\pi(\theta, \omega) \in \Delta(\Theta \times \Omega)$

or

Mechanism and Bidder Utility

Definition: Mechanism

A (direct revelation) mechanism, $(\boldsymbol{p}, \boldsymbol{x})$, is defined by, given the bidder type and external signal (θ, ω), the probability that the seller allocates the item to the bidder, $\boldsymbol{p}(\theta, \omega)$, and a monetary transfer from the bidder to the seller, $\boldsymbol{x}(\theta, \omega)$.

Mechanism and Bidder Utility

Definition: Mechanism

A (direct revelation) mechanism, $(\boldsymbol{p}, \boldsymbol{x})$, is defined by, given the bidder type and external signal (θ, ω), the probability that the seller allocates the item to the bidder, $\boldsymbol{p}(\theta, \omega)$, and a monetary transfer from the bidder to the seller, $\boldsymbol{x}(\theta, \omega)$.

Definition: Bidder Utility

Given a realization of the external signal ω, reported type $\theta' \in \Theta$ by the bidder, and true type $\theta \in \Theta$, the bidder's utility under mechanism $(\boldsymbol{p}, \boldsymbol{x})$ is:

$$U(\theta, \theta', \omega) = v(\theta)\boldsymbol{p}(\theta', \omega) - \boldsymbol{x}(\theta', \omega)$$

Definition: Ex-Interim Individual Rationality (IR)

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

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

Definition: Ex-Interim Individual Rationality (IR)

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

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

Definition: Bayesian Incentive Compatibility (IC)

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

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

Definition: Optimal Mechanisms

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

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

Full Surplus Extraction with Bayesian Mechanisms (Cremer and McLean 1985; A, Conitzer, and Lopomo 2016)

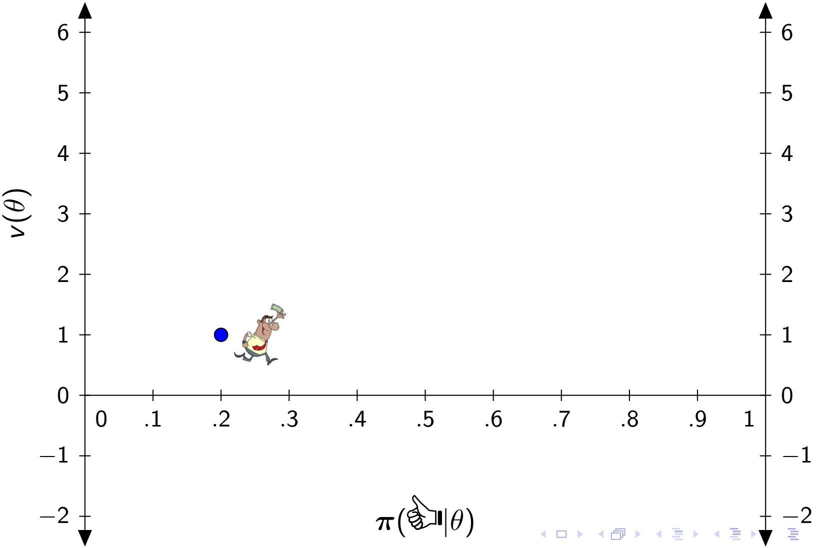

Full Surplus Extraction with Bayesian Mechanisms (Cremer and McLean 1985; A, Conitzer, and Lopomo 2016)

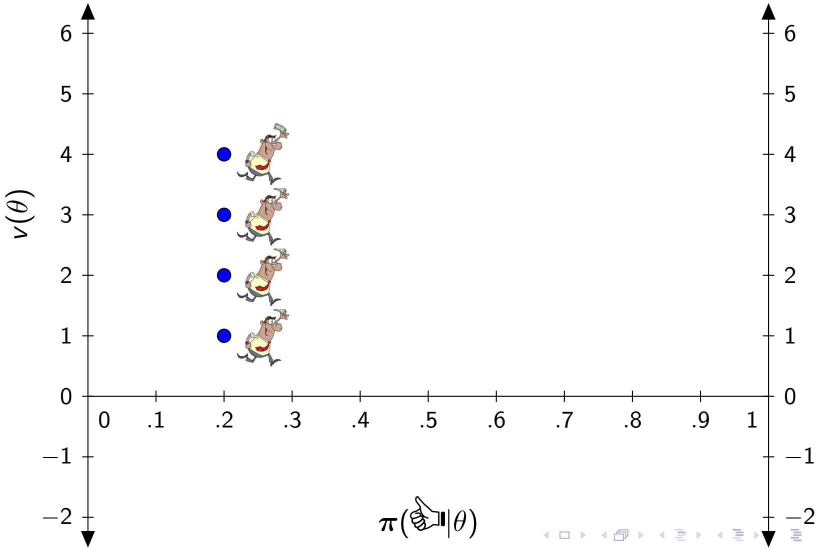

Full Surplus Extraction with Bayesian Mechanisms (Cremer and McLean 1985; A, Conitzer, and Lopomo 2016)

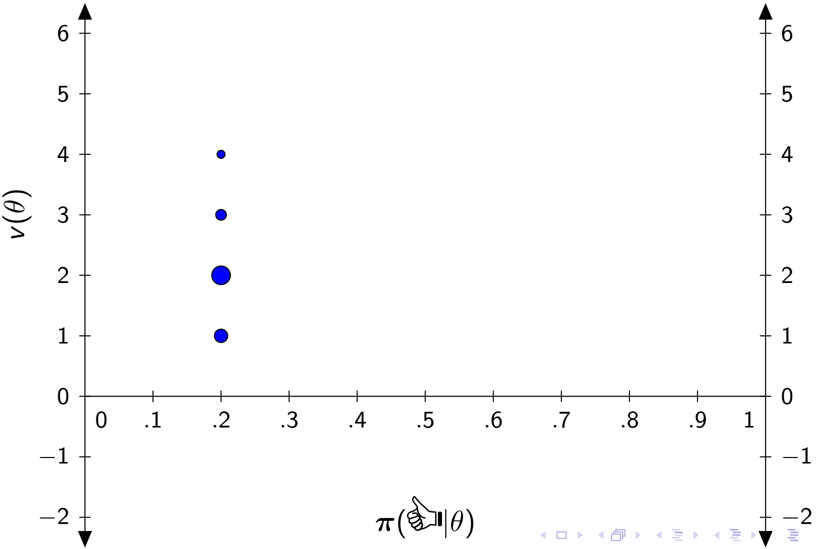

Full Surplus Extraction with Bayesian Mechanisms (Cremer and McLean 1985; A, Conitzer, and Lopomo 2016)

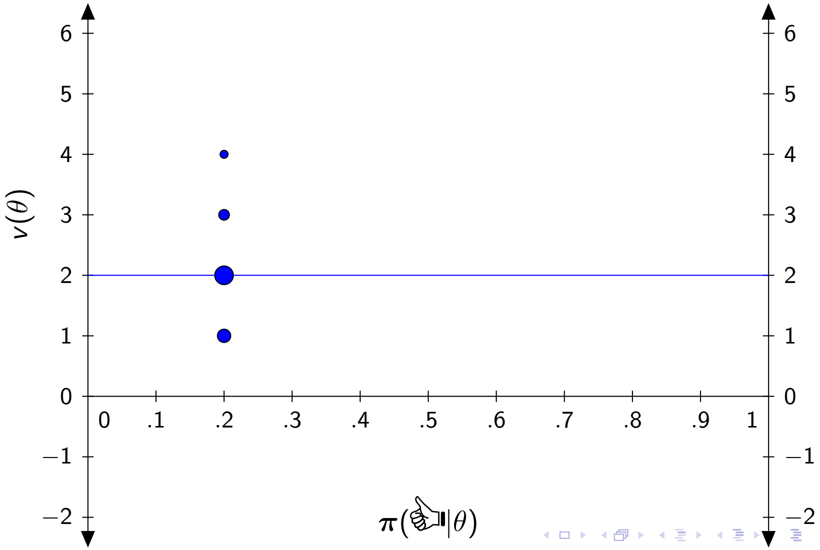

Full Surplus Extraction with Bayesian Mechanisms (Cremer and McLean 1985; A, Conitzer, and Lopomo 2016)

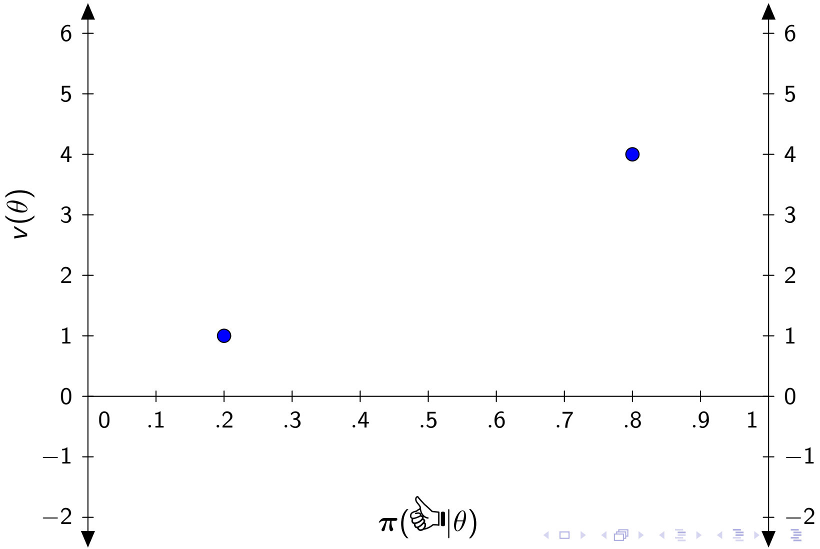

Full Surplus Extraction with Bayesian Mechanisms (Cremer and McLean 1985; A, Conitzer, and Lopomo 2016)

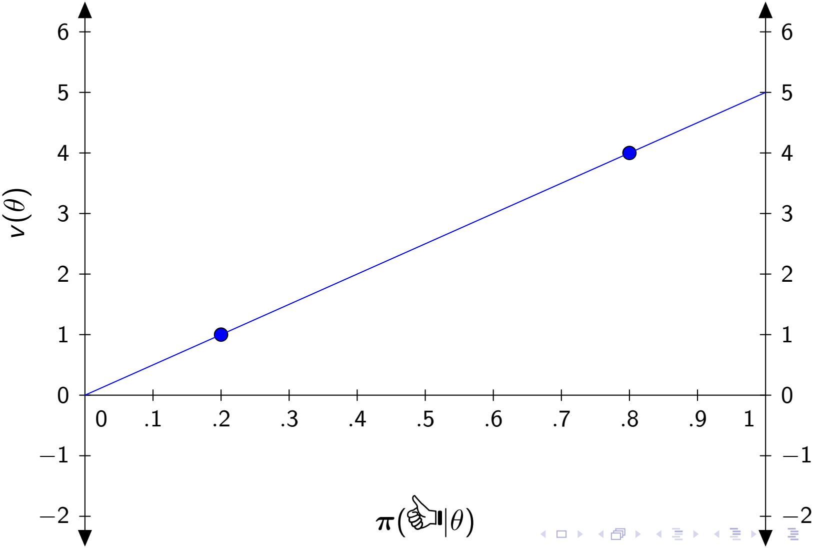

Full Surplus Extraction with Bayesian Mechanisms (Cremer and McLean 1985; A, Conitzer, and Lopomo 2016)

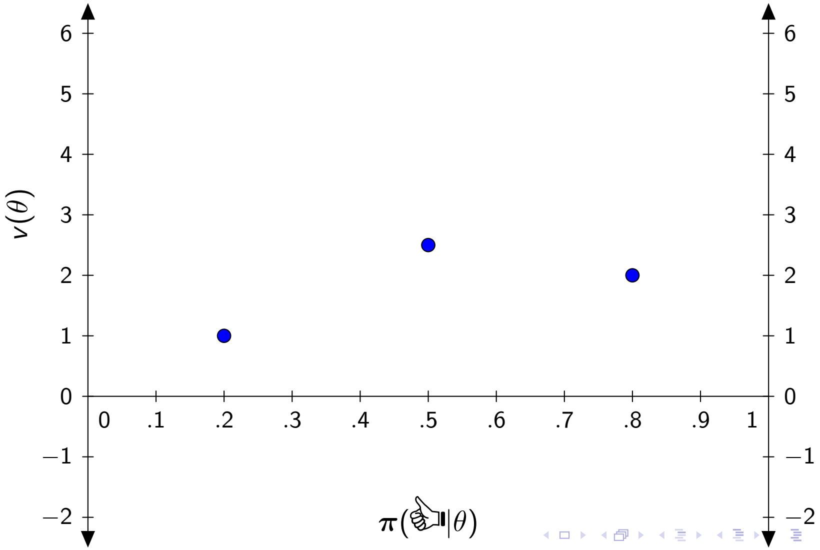

Full Surplus Extraction with Bayesian Mechanisms (Cremer and McLean 1985; A, Conitzer, and Lopomo 2016)

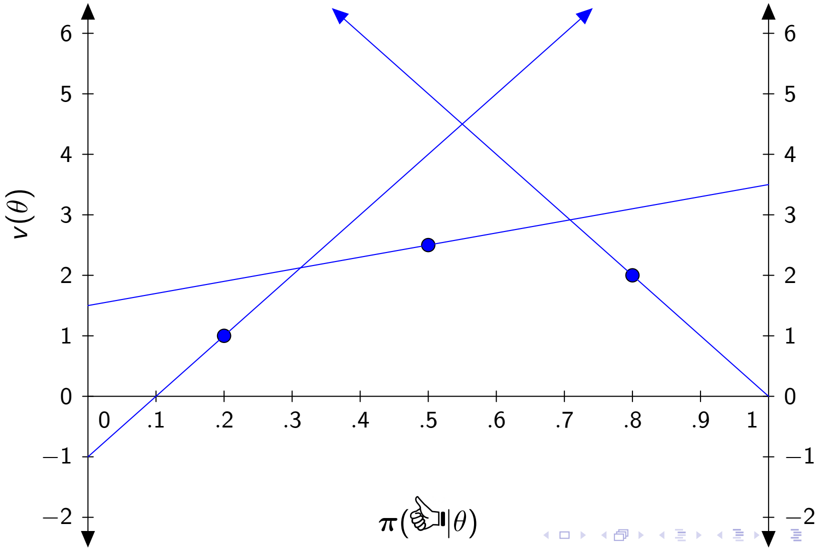

Full Surplus Extraction with Bayesian Mechanisms (Cremer and McLean 1985; A, Conitzer, and Lopomo 2016)

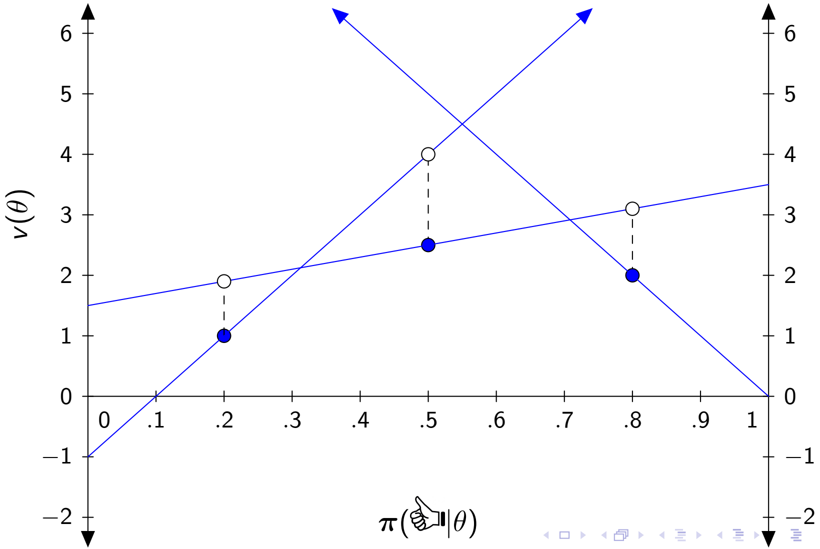

Full Surplus Extraction with Bayesian Mechanisms (Cremer and McLean 1985; A, Conitzer, and Lopomo 2016)

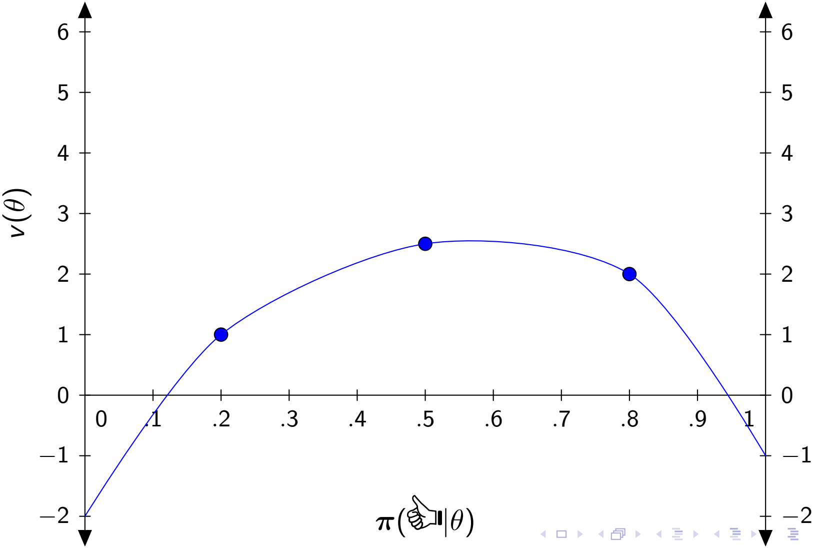

Full Surplus Extraction with Bayesian Mechanisms (Cremer and McLean 1985; A, Conitzer, and Lopomo 2016)

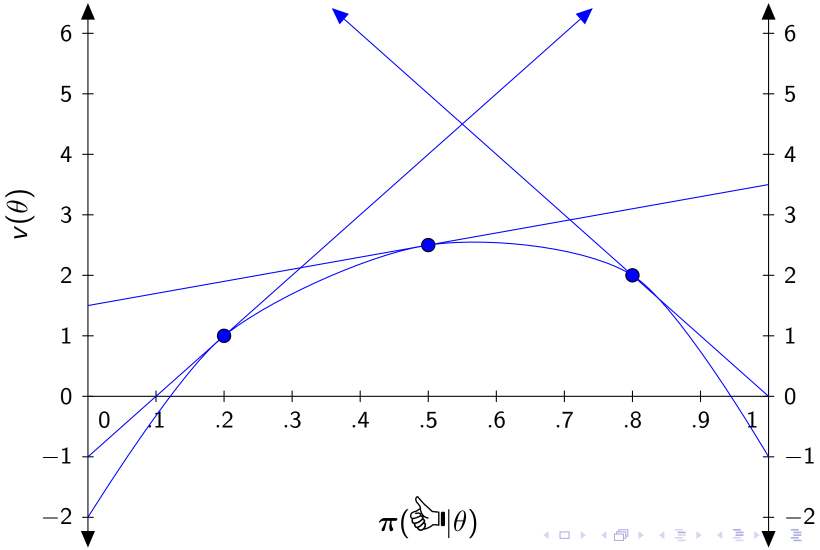

Uncertain Distributions

- What if we don't know the true distribution?
    - Maybe we observe samples from previous auction rounds
- Full extraction is still possible and easy with a finite set of potential distributions
    - Lopomo, Rigotti, and Shannon 2009 give conditions under which full extraction is possible with Knightian uncertainty in a discrete type space
    - Fu et. al. 2014 find that a single sample from the underlying distribution is sufficient to extract full revenue (given a generic condition)
- We look at an infinite set of distributions
    - Discrete set for impossibility result
    - Single bidder and external signal, bidder knows true distribution
    - We know the marginal distribution over bidder types
    - Finite number of samples from the true distribution
    - Bidders report both type and true distribution

Uncertain Distributions

- What if we don't know the true distribution?
  - Maybe we observe samples from previous auction rounds
- Full extraction is still possible and easy with a finite set of potential distributions
  - Lopomo, Rigotti, and Shannon 2009 give conditions under which full extraction is possible with Knightian uncertainty in a discrete type space
  - Fu et. al. 2014 find that a single sample from the underlying distribution is sufficient to extract full revenue (given a generic condition)
- We look at an infinite set of distributions
  - Discrete set for impossibility result
  - Single bidder and external signal, bidder knows true distribution
  - We know the marginal distribution over bidder types
  - Finite number of samples from the true distribution
  - Bidders report both type and true distribution

Uncertain Distributions

- What if we don't know the true distribution?
  - Maybe we observe samples from previous auction rounds
- Full extraction is still possible and easy with a finite set of potential distributions
  - Lopomo, Rigotti, and Shannon 2009 give conditions under which full extraction is possible with Knightian uncertainty in a discrete type space
  - Fu et. al. 2014 find that a single sample from the underlying distribution is sufficient to extract full revenue (given a generic condition)
- We look at an infinite set of distributions
  - Discrete set for impossibility result
  - Single bidder and external signal, bidder knows true distribution
  - We know the marginal distribution over bidder types
  - Finite number of samples from the true distribution
  - Bidders report both type and true distribution

Uncertain Distributions

- What if we don't know the true distribution?
  - Maybe we observe samples from previous auction rounds
- Full extraction is still possible and easy with a finite set of potential distributions
  - Lopomo, Rigotti, and Shannon 2009 give conditions under which full extraction is possible with Knightian uncertainty in a discrete type space
  - Fu et. al. 2014 find that a single sample from the underlying distribution is sufficient to extract full revenue (given a generic condition)
- We look at an infinite set of distributions
  - Discrete set for impossibility result
  - Single bidder and external signal, bidder knows true distribution
  - We know the marginal distribution over bidder types
  - Finite number of samples from the true distribution
  - Bidders report both type and true distribution

Converging Sequences of Distributions

Definition: Converging Distributions

A countably infinite sequence of distributions $\{\boldsymbol{\pi}_i\}_{i=1}^{\infty}$ is said to be **converging to the distribution** $\boldsymbol{\pi}^*$, the **convergence point**, if for all $\theta \in \Theta$ and $\epsilon > 0$, there exists a $T \in \mathbb{N}$ such that for all $i \geq T$, $||\boldsymbol{\pi}_i(\cdot|\theta) - \boldsymbol{\pi}^*(\cdot|\theta)|| < \epsilon$. I.e., for each $\theta \in \Theta$, the conditional distributions in the sequence, $\{\boldsymbol{\pi}_i(\cdot|\theta)\}_{i=1}^{\infty}$, converge to the conditional distribution $\boldsymbol{\pi}^*(\cdot|\theta)$ in the l^2 norm.

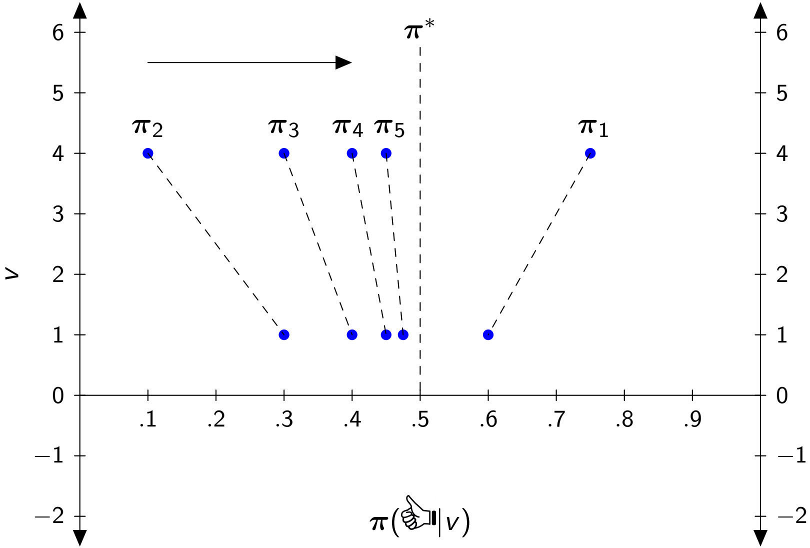

π^*
π_2
π_3
π_4
π_5
π_1
v
6
5
4
3
2
1
0
−1
−2
.1 .2 .3 .4 .5 .6 .7 .8 .9
$\pi(\cdot|v)$

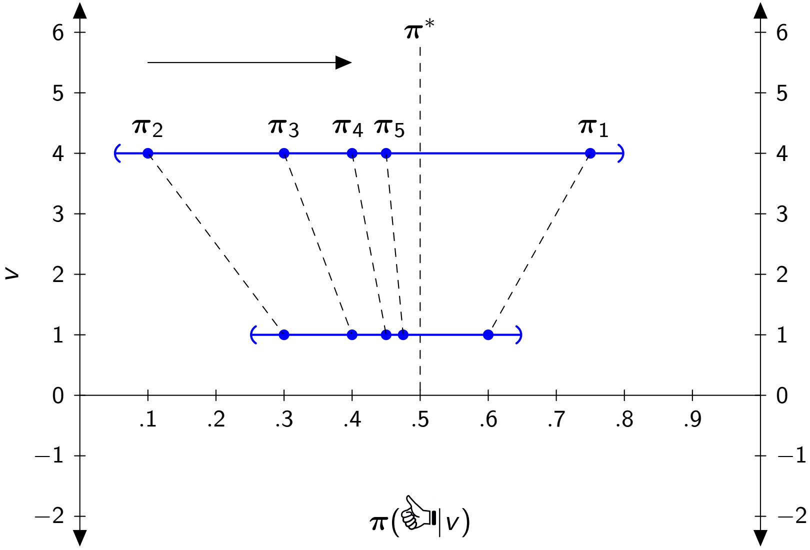

π^*
π_2
π_3
π_4
π_5
π_1
6
5
4
3
2
1
0
−1
−2
v
.1
.2
.3
.4
.5
.6
.7
.8
.9
$\pi(👍|v)$

Distribution as Private Information

Definition: Mechanism with Private Distributions

A (direct revelation) mechanism, $(\boldsymbol{p}, \boldsymbol{x})$, is defined by, given a bidder type, a distribution, and the external signal, $(\theta, \boldsymbol{\pi}, \omega)$, the probability that the seller allocates the item to the bidder, $\boldsymbol{p}(\theta, \boldsymbol{\pi}, \omega)$, and a monetary transfer from the bidder to the seller, $\boldsymbol{x}(\theta, \boldsymbol{\pi}, \omega)$.

Distribution as Private Information

Definition: Mechanism with Private Distributions

A (direct revelation) mechanism, $(\boldsymbol{p}, \boldsymbol{x})$, is defined by, given a bidder type, a distribution, and the external signal, $(\theta, \boldsymbol{\pi}, \omega)$, the probability that the seller allocates the item to the bidder, $\boldsymbol{p}(\theta, \boldsymbol{\pi}, \omega)$, and a monetary transfer from the bidder to the seller, $\boldsymbol{x}(\theta, \boldsymbol{\pi}, \omega)$.

Definition: Bidder Utility with Private Distributions

Given a realization of the external signal ω, reported type $\theta' \in \Theta$ by the bidder, reported distribution $\boldsymbol{\pi}' \in \{\boldsymbol{\pi}_i\}_{i=1}^{\infty}$, true type $\theta \in \Theta$, and true distribution $\boldsymbol{\pi} \in \{\boldsymbol{\pi}_i\}_{i=1}^{\infty}$, the bidder's utility under mechanism $(\boldsymbol{p}, \boldsymbol{x})$ is:

$$U(\theta, \boldsymbol{\pi}, \theta', \boldsymbol{\pi}', \omega) = v(\theta)\boldsymbol{p}(\theta', \boldsymbol{\pi}', \omega) - \boldsymbol{x}(\theta', \boldsymbol{\pi}', \omega)$$

Definition: Ex-Interim Individual Rationality (IR)

A mechanism $(\boldsymbol{p}, \boldsymbol{x})$ is *ex-interim individually rational (IR)* if for all $\theta \in \Theta$ and $\boldsymbol{\pi} \in \{\boldsymbol{\pi}_i\}_{i=1}^{\infty}$:

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

Definition: Ex-Interim Individual Rationality (IR)

A mechanism $(\boldsymbol{p}, \boldsymbol{x})$ is *ex-interim individually rational (IR)* if for all $\theta \in \Theta$ and $\boldsymbol{\pi} \in \{\boldsymbol{\pi}_i\}_{i=1}^{\infty}$:

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

Definition: Bayesian Incentive Compatibility (IC)

A mechanism $(\boldsymbol{p}, \boldsymbol{x})$ is *Bayesian incentive compatible (IC)* if for all $\theta, \theta' \in \Theta$ and $\boldsymbol{\pi}, \boldsymbol{\pi}' \in \{\boldsymbol{\pi}_i\}_{i=1}^{\infty}$:

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

Convergence to an Interior Point

Assumption: Converging to an Interior Point

For the sequence of distributions $\{\boldsymbol{\pi}_i\}_{i=1}^{\infty}$ *converging to* $\boldsymbol{\pi}^*$ *and for any* $\theta' \in \Theta$, *there exists a subset of distributions of size* $|\Omega|$ *from the set* $\{\boldsymbol{\pi}_i(\cdot|\theta)\}_{i,\theta}$ *that is affinely independent and the distribution* $\boldsymbol{\pi}^*(\cdot|\theta')$ *is a strictly convex combination of the elements of the subset. I.e., there exists* $\{\alpha_k\}_{k=1}^{|\Omega|}$, $\alpha_k \in (0,1)$, *and* $\{\boldsymbol{\pi}_k(\cdot|\theta_k)\}_{k=1}^{|\Omega|}$ *such that* $\boldsymbol{\pi}^*(\cdot|\theta') = \sum_{k=1}^{|\Omega|} \alpha_k \boldsymbol{\pi}_k(\cdot|\theta_k)$.

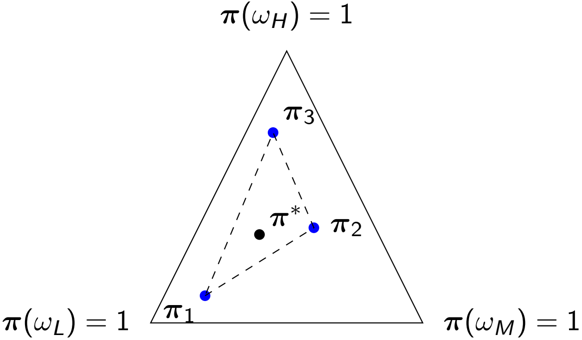
$\pi(\omega_H) = 1$
π_3
π^*
π_2
π_1
$\pi(\omega_L) = 1$
$\pi(\omega_M) = 1$

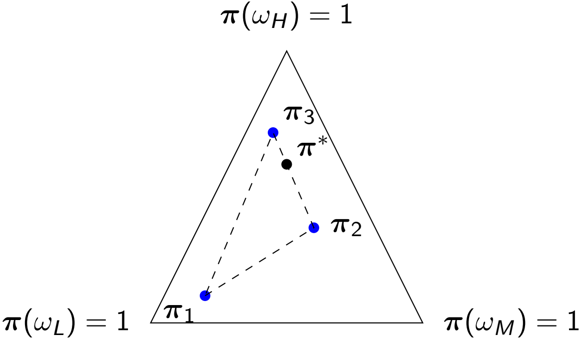
$\pi(\omega_H) = 1$
π_3
π^*
π_2
π_1
$\pi(\omega_L) = 1$
$\pi(\omega_M) = 1$

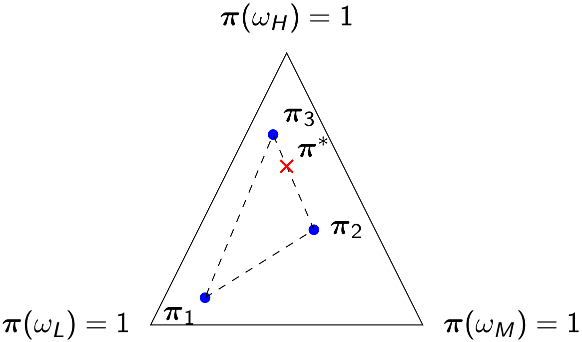
$\pi(\omega_H) = 1$
π_3
π^*
π_2
π_1
$\pi(\omega_L) = 1$
$\pi(\omega_M) = 1$

Inapproximability of the Optimal Mechanism

Theorem: Inapproximability of the Optimal Mechanism

Let $\{\boldsymbol{\pi}_i\}_{i=1}^{\infty}$ be a sequence of distributions converging to $\boldsymbol{\pi}^*$. Denote the revenue of the optimal mechanism for the distribution $\boldsymbol{\pi}^*$ by R. For any $k > 0$, there exists a $T \in \mathbb{N}$ such that for all $\boldsymbol{\pi}_{i'} \in \{\boldsymbol{\pi}_i\}_{i=T}^{\infty}$, the expected revenue is less than $R + k$.

Inapproximability of the Optimal Mechanism

Theorem: Inapproximability of the Optimal Mechanism

Let $\{\boldsymbol{\pi}_i\}_{i=1}^{\infty}$ be a sequence of distributions converging to $\boldsymbol{\pi}^*$. Denote the revenue of the optimal mechanism for the distribution $\boldsymbol{\pi}^*$ by R. For any $k > 0$, there exists a $T \in \mathbb{N}$ such that for all $\boldsymbol{\pi}_{i'} \in \{\boldsymbol{\pi}_i\}_{i=T}^{\infty}$, the expected revenue is less than $R + k$.

Corrollary: Sampling Doesn't Help

The above still holds if the mechanism designer has access to a finite number of samples from the underlying true distribution.

Sufficient Correlation Implies Near Optimal Revenue

Theorem: Sufficient Correlation Implies Near Optimal Revenue

For any distribution $\boldsymbol{\pi}^*$ that satisfies the ACL condition with optimal revenue R and given any positive constant $k > 0$, there exists $\epsilon > 0$ and a mechanism such that for all distributions, $\boldsymbol{\pi}'$, for which for all $\theta \in \Theta$, $||\boldsymbol{\pi}^*(\cdot|\theta) - \boldsymbol{\pi}'(\cdot|\theta)|| < \epsilon$, the revenue generated by the mechanism is greater than or equal to $R - k$.

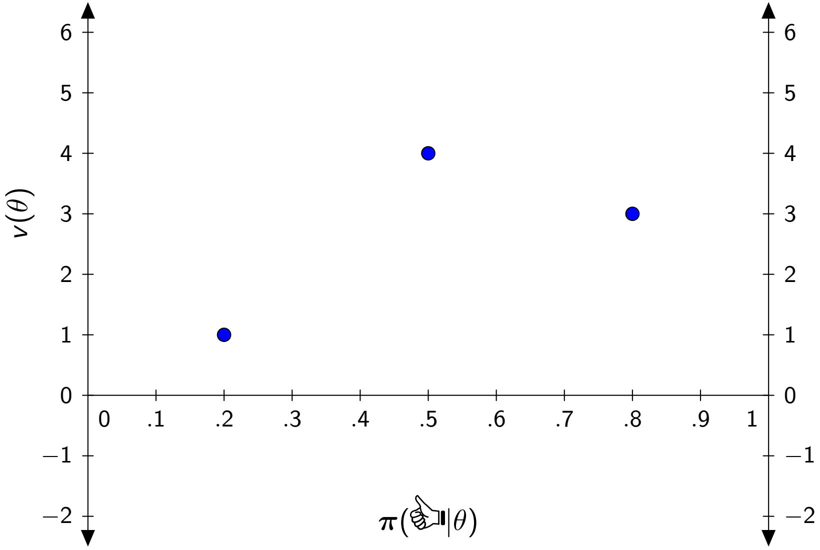
6
5
4
3
2
1
0
−1
−2
$v(\theta)$
0 .1 .2 .3 .4 .5 .6 .7 .8 .9 1
$\pi(\cdot|\theta)$

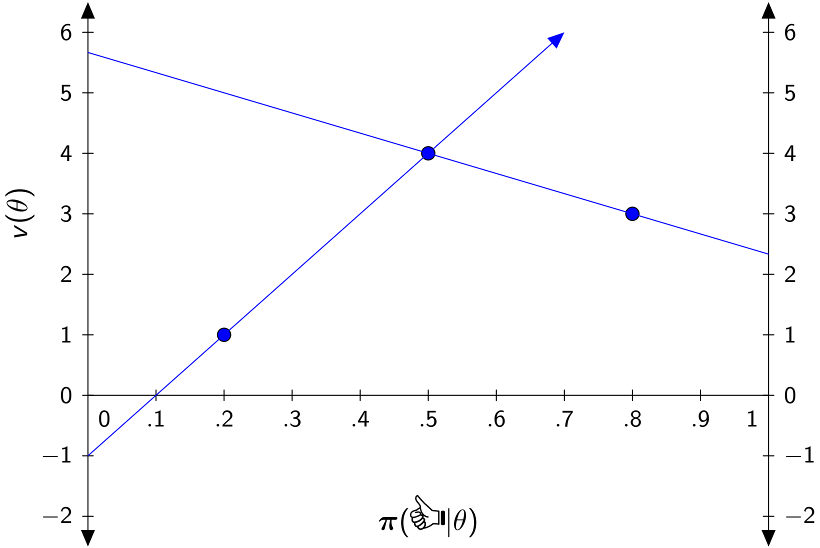
6
5
4
3
2
1
0
−1
−2
$v(\theta)$
0 .1 .2 .3 .4 .5 .6 .7 .8 .9 1
$\pi($👍$|\theta)$

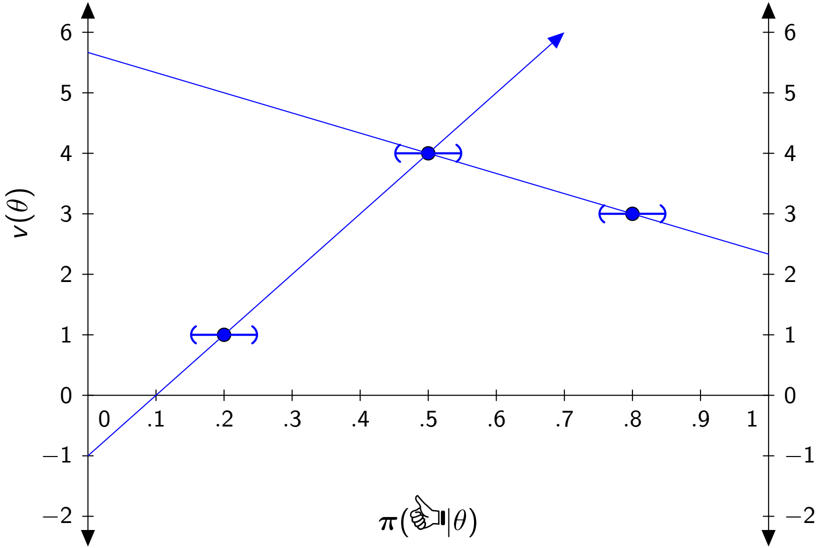
6
5
4
3
2
1
0
−1
−2
$v(\theta)$
0
.1
.2
.3
.4
.5
.6
.7
.8
.9
1
$\pi($👍$|\theta)$

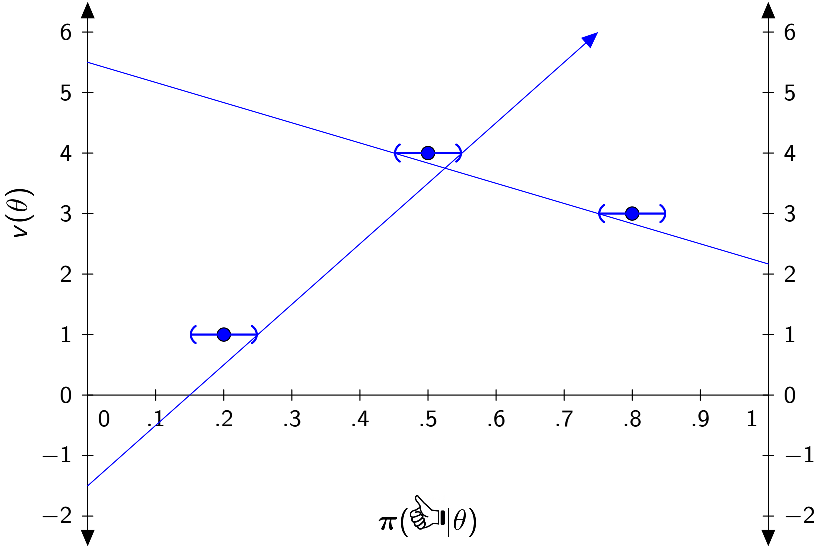
6
5
4
3
2
1
0
−1
−2
$v(\theta)$
0
.1
.2
.3
.4
.5
.6
.7
.8
.9
1
$\pi(\cdot|\theta)$

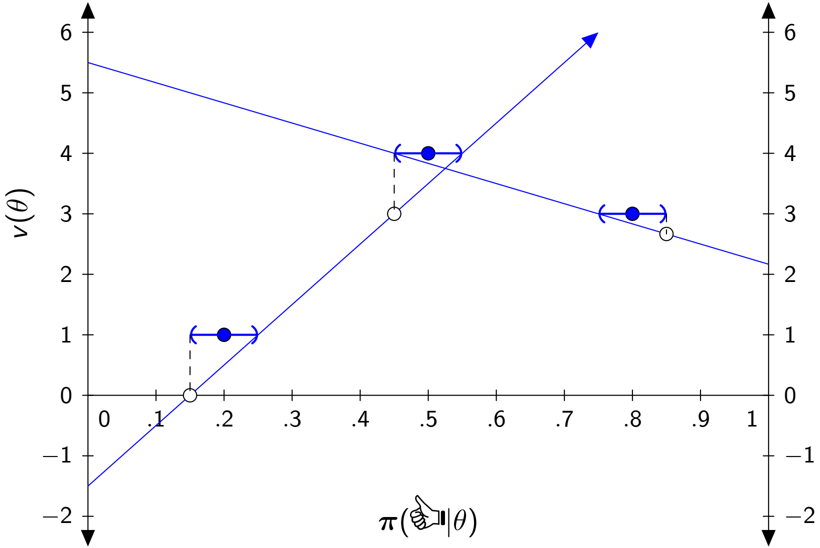
6
5
4
3
2
1
0
−1
−2
0
.1
.2
.3
.4
.5
.6
.7
.8
.9
1
$v(\theta)$
$\pi(☞|\theta)$

A, Conitzer, and Stone 2017 - AAAI - Automated Design of Robust Mechanisms

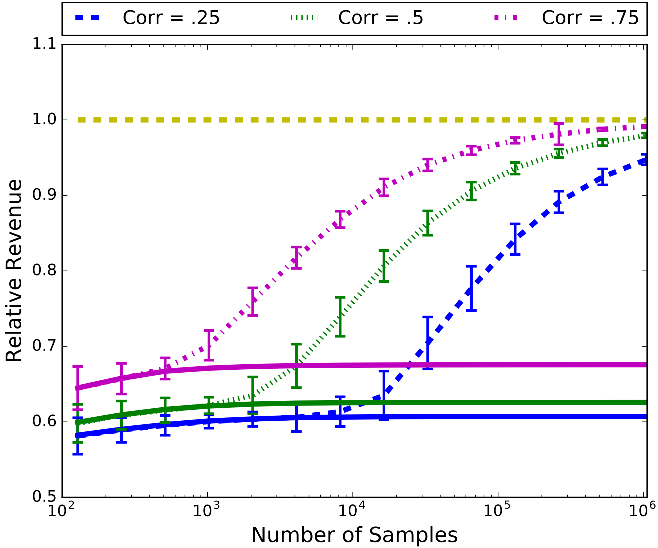

Related Work

- Unknown Correlated Distributions (Lopomo, Rigotti, and Shannon 2009, Fu, Haghpanah, Hartline, and Kleinberg 2014)
- 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 Distributions (Elkind 2007, 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?

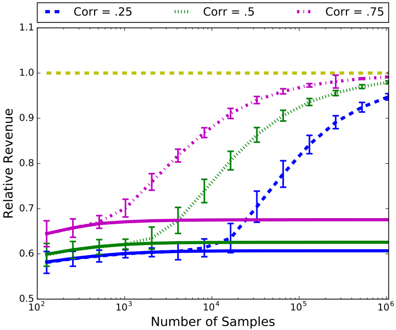

I will also be presenting this as a poster at DD-2 during the Thursday morning poster session. Please come by!