Automated Mechanism Design for Correlated Valuations

Vincent Conitzer; joint work with:

Michael Albert (Duke → UVA)

Giuseppe (Pino) Lopomo (Duke)

Peter Stone (UT Austin)

Paper:

Mechanism Design for Correlated Valuations: Efficient Methods for Revenue Maximization. **Michael Albert**, Vincent Conitzer, Giuseppe Lopomo, and Peter Stone. *Operations Research,* forthcoming.

Automated mechanism design input

Instance is given by

Set of possible *outcomes*

Set of *agents*

For each agent

set of possible *types*

probability distribution over these types

Objective function

Gives a value for each outcome for each combination of agents' types

E.g., social welfare, revenue

Restrictions on the mechanism

Are payments allowed?

Is randomization over outcomes allowed?

What versions of incentive compatibility (IC) & individual rationality (IR) are used?

How hard is designing an optimal *deterministic* mechanism (without reporting costs)?

[C. & Sandholm UAI'02, ICEC'03, EC'04]

NP-complete (even with 1 reporting agent):	**Solvable in polynomial time (for any *constant* number of agents):**
1. Maximizing social welfare (no payments) 2. Designer's own utility over outcomes (no payments) 3. General (linear) objective that doesn't regard payments 4. Expected revenue	1. Maximizing social welfare (not regarding the payments) (VCG)

1 and 3 hold even with no IR constraints

Positive results (randomized mechanisms)

[C. & Sandholm UAI'02, ICEC'03, EC'04]

- Use linear programming
- Variables:

  $p(o \mid \theta_1, \ldots, \theta_n)$ = probability that outcome o is chosen given types $\theta_1, \ldots, \theta_n$

  (maybe) $\pi_i(\theta_1, \ldots, \theta_n)$ = i's payment given types $\theta_1, \ldots, \theta_n$
- Strategy-proofness constraints: for all $i, \theta_1, \ldots \theta_n, \theta_i'$:

  $\Sigma_o p(o \mid \theta_1, \ldots, \theta_n) u_i(\theta_i, o) + \pi_i(\theta_1, \ldots, \theta_n) \geq$

  $\Sigma_o p(o \mid \theta_1, \ldots, \theta_i', \ldots, \theta_n) u_i(\theta_i, o) + \pi_i(\theta_1, \ldots, \theta_i', \ldots, \theta_n)$
- Individual-rationality constraints: for all $i, \theta_1, \ldots \theta_n$:

  $\Sigma_o p(o \mid \theta_1, \ldots, \theta_n) u_i(\theta_i, o) + \pi_i(\theta_1, \ldots, \theta_n) \geq 0$
- Objective (e.g., sum of utilities)

  $\Sigma_{\theta_1, \ldots, \theta_n} p(\theta_1, \ldots, \theta_n) \Sigma_i (\Sigma_o p(o \mid \theta_1, \ldots, \theta_n) u_i(\theta_i, o) + \pi_i(\theta_1, \ldots, \theta_n))$
- Also works for BNE incentive compatibility, ex-interim individual rationality notions, other objectives, etc.
- For deterministic mechanisms, can still use mixed integer programming: require probabilities in $\{0, 1\}$
  - Remember typically designing the optimal deterministic mechanism is NP-hard

A simple example

One item for sale (free disposal)
2 agents, IID valuations: uniform over {1, 2}
Maximize expected revenue under ex-interim IR, Bayes-Nash equilibrium
How much can we get?
(What is optimal expected welfare?)

Agent 1's valuation \ Agent 2's valuation	1	2
1	0.25	0.25
2	0.25	0.25

probabilities

A simple example

One item for sale (free disposal)

2 agents, IID valuations: uniform over {1, 2}

Maximize expected revenue under ex-interim IR, Bayes-Nash equilibrium

How much can we get?

(What is optimal expected welfare?)

Agent 2's valuation

Agent 1's valuation	1	2
1	0.25	0.25
2	0.25	0.25

probabilities

Our old AMD solver [C. & Sandholm, 2002, 2003] gives:

```
Status:      OPTIMAL
Objective:   obj = 1.5 (MAXimum)
```

[nonzero variables:]

`p_t_1_1_o3`	1	(probability of disposal for (1, 1))
`p_t_2_1_o1`	1	(probability 1 gets the item for (2, 1))
`p_t_1_2_o2`	1	(probability 2 gets the item for (1, 2))
`p_t_2_2_o2`	1	(probability 2 gets the item for (2, 2))
`pi_2_2_1`	2	(1's payment for (2, 2))
`pi_2_2_2`	4	(2's payment for (2, 2))

A slightly different distribution

One item for sale (free disposal)
2 agents, valuations drawn as on right
Maximize expected revenue under ex-interim IR, Bayes-Nash equilibrium
How much can we get?
(What is optimal expected welfare?)

Agent 2's valuation (columns), *Agent 1's valuation* (rows)

	1	2
1	0.251	0.250
2	0.250	0.249

probabilities

A slightly different distribution

One item for sale (free disposal)

2 agents, valuations drawn as on right

Maximize expected revenue under ex-interim IR, Bayes-Nash equilibrium

How much can we get?

(What is optimal expected welfare?)

Agent 2's valuation (columns), *Agent 1's valuation* (rows)

	1	2
1	0.251	0.250
2	0.250	0.249

probabilities

```
Status:       OPTIMAL
Objective:    obj = 1.749 (MAXimum)
```

[*some of the nonzero payment variables:*]

```
pi_1_1_2          62501
pi_2_1_2         -62750
pi_2_1_1          2
pi_1_2_2          3.992
```

You'd better be really sure about your distribution!

A nearby distribution without correlation

One item for sale (free disposal)

2 agents, valuations IID: 1 w/ .501, 2 w/ .499

Maximize expected revenue under ex-interim IR, Bayes-Nash equilibrium

How much can we get?

(What is optimal expected welfare?)

Agent 2's valuation (columns); *Agent 1's valuation* (rows)

	1	2
1	0.251001	0.249999
2	0.249999	0.249001

probabilities

A nearby distribution without correlation

One item for sale (free disposal)

2 agents, valuations IID: 1 w/ .501, 2 w/ .499

Maximize expected revenue under ex-interim IR, Bayes-Nash equilibrium

How much can we get?

(What is optimal expected welfare?)

Agent 1's valuation \ Agent 2's valuation	1	2
1	0.251001	0.249999
2	0.249999	0.249001

probabilities

```
Status:       OPTIMAL
Objective:    obj = 1.499 (MAXimum)
```

Cremer-McLean [1985]

For every agent, consider the following matrix Γ of conditional probabilities, where Θ is the set of types for the agent and Ω is the set of signals (joint types for other agents, or something else observable to the auctioneer)

$$\mathbf{\Gamma} = \begin{bmatrix} \pi(1|1) & \cdots & \pi(|\Omega||1) \\ \vdots & \ddots & \vdots \\ \pi(1||\Theta|) & \cdots & \pi(|\Omega|||\Theta|) \end{bmatrix}$$

If Γ has rank |Θ| for every agent then the auctioneer can allocate efficiently and extract the full surplus as revenue (!!)

Preview of Results

Automated mechanism design procedure that is both *computationally efficient* and *sample efficient* for **maximizing revenue** when there is *sufficient* correlation.

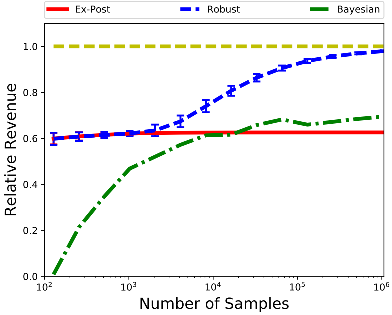

Outline

- **Known Distributions** (AAAI 2015; 2016)
- **Robust Mechanism Design for Revenue Maximizing Mechanisms** (AAMAS 2017, AAAI 2017, to appear in OR 2021)

Outline

- **Known Distributions** (AAAI 2015; 2016)
- **Robust Mechanism Design for Revenue Maximizing Mechanisms** (AAMAS 2017, AAAI 2017, to appear in OR 2021)

Outline

- **Known Distributions** (AAAI 2015; 2016)
- **Robust Mechanism Design for Revenue Maximizing Mechanisms** (AAMAS 2017, AAAI 2017, to appear in OR 2021)

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

Mechanism Design with Correlated Distributions

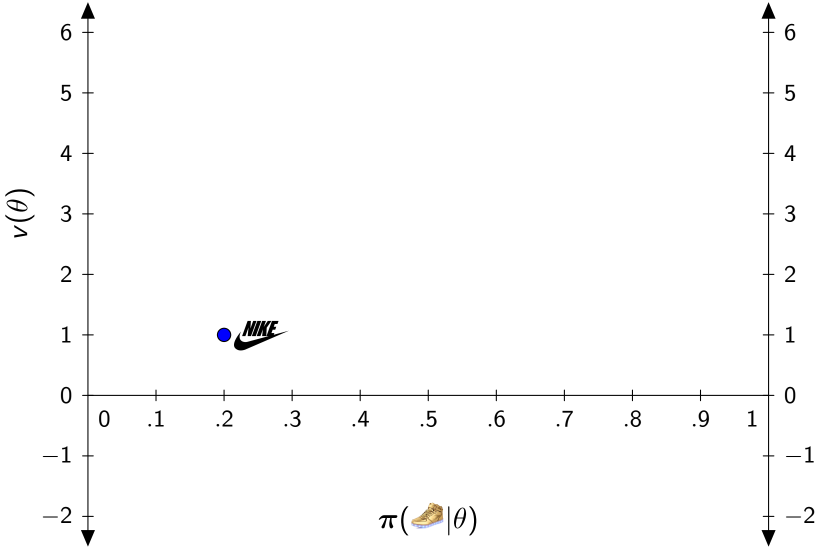

Mechanism Design with Correlated Distributions

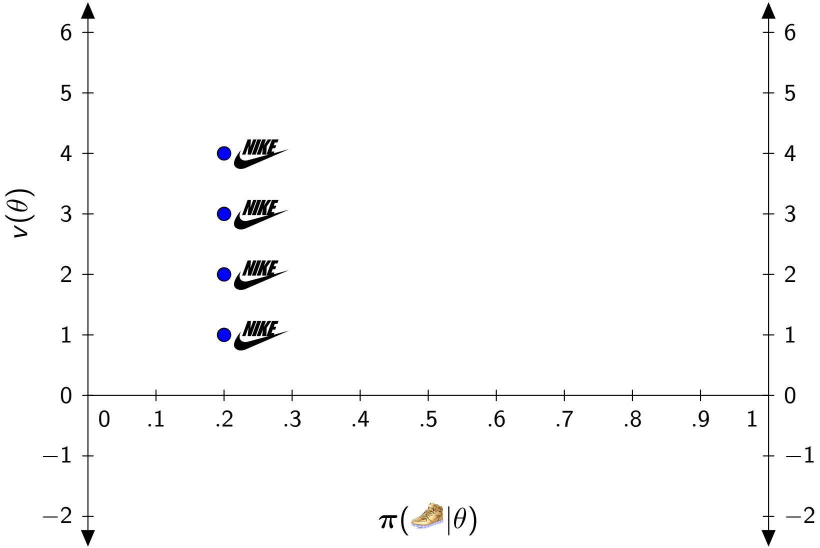

Mechanism Design with Correlated Distributions

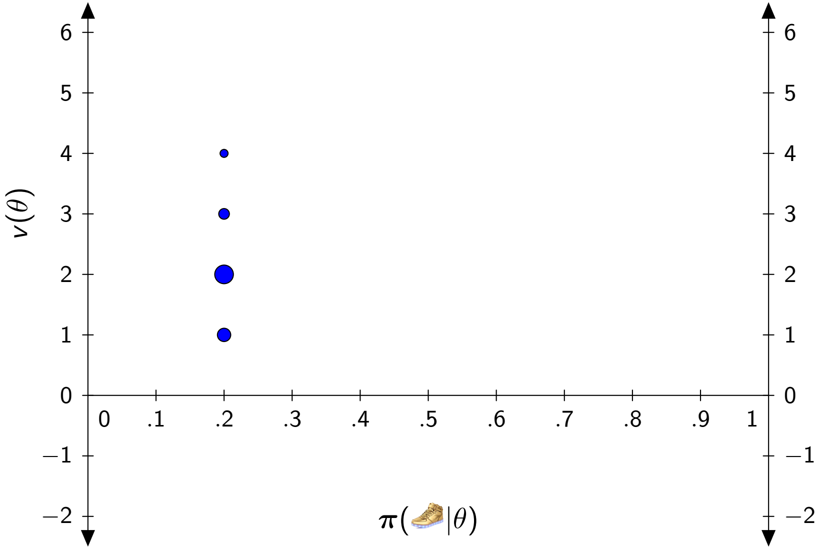

Mechanism Design with Correlated Distributions

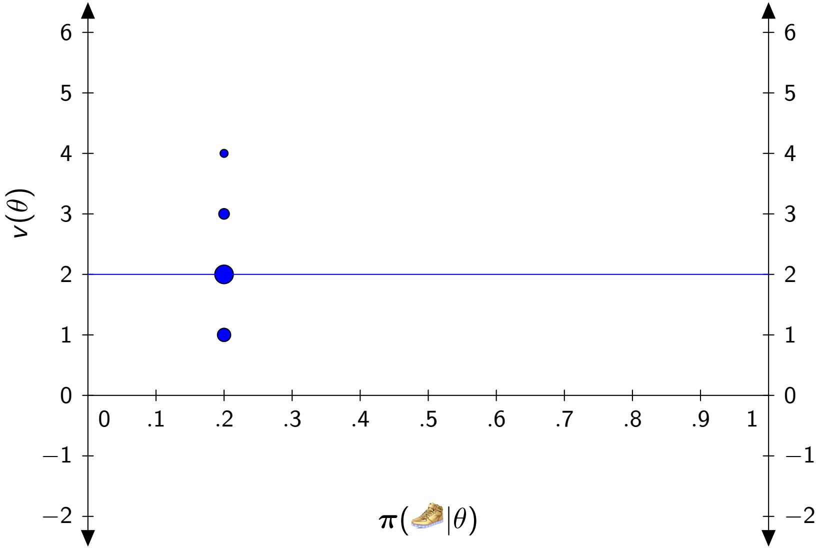

Mechanism Design with Correlated Distributions

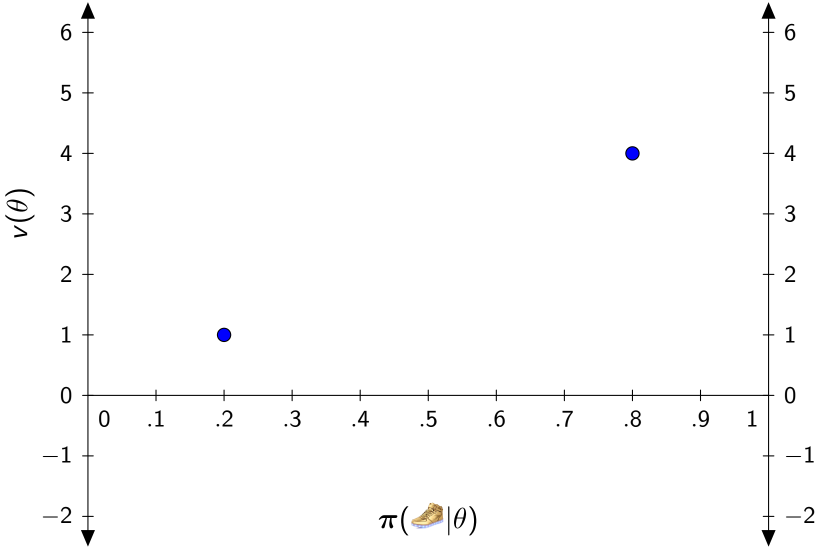

Mechanism Design with Correlated Distributions

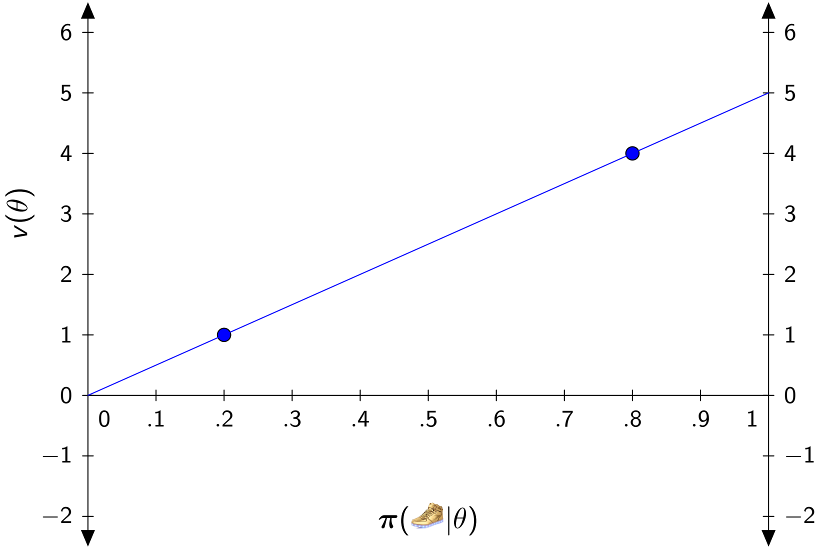

Mechanism Design with Correlated Distributions

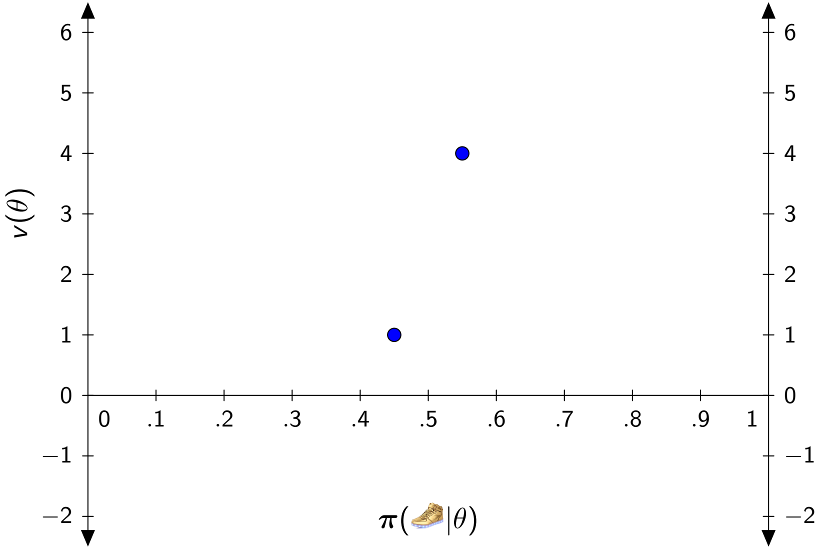

Mechanism Design with Correlated Distributions

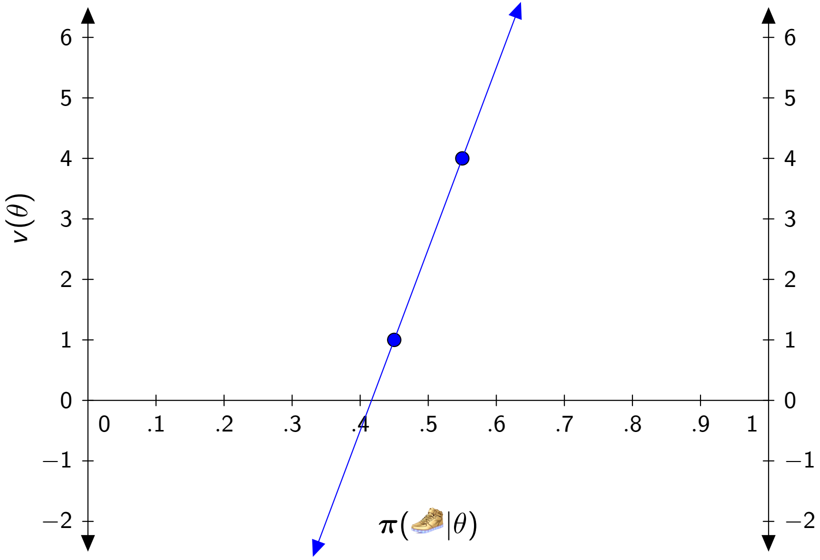

Necessary and Sufficient Condition

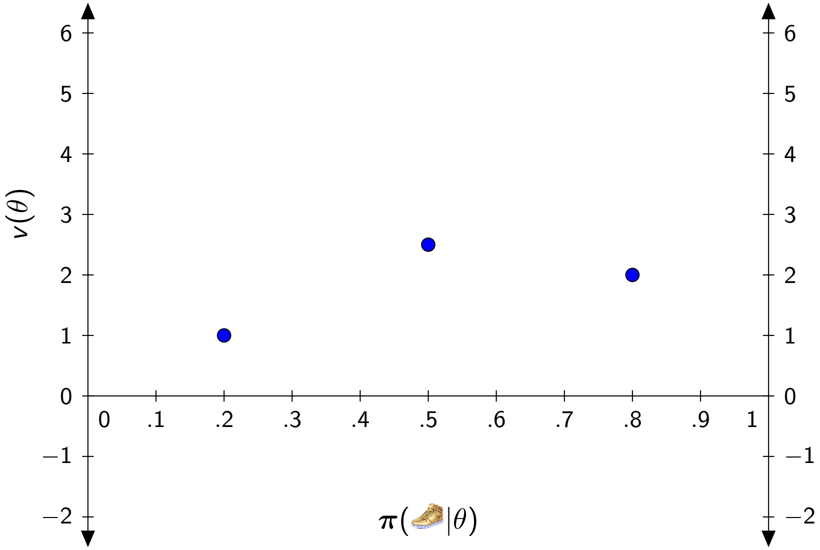

Necessary and Sufficient Condition

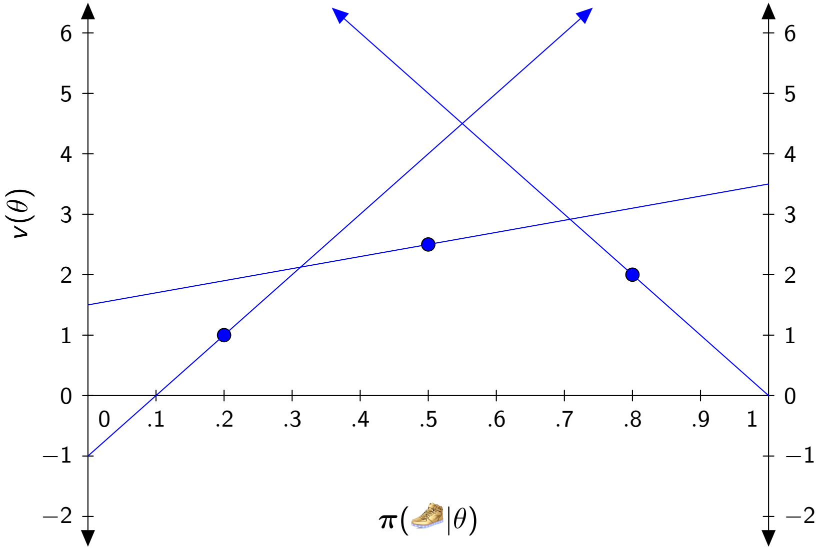

Necessary and Sufficient Condition

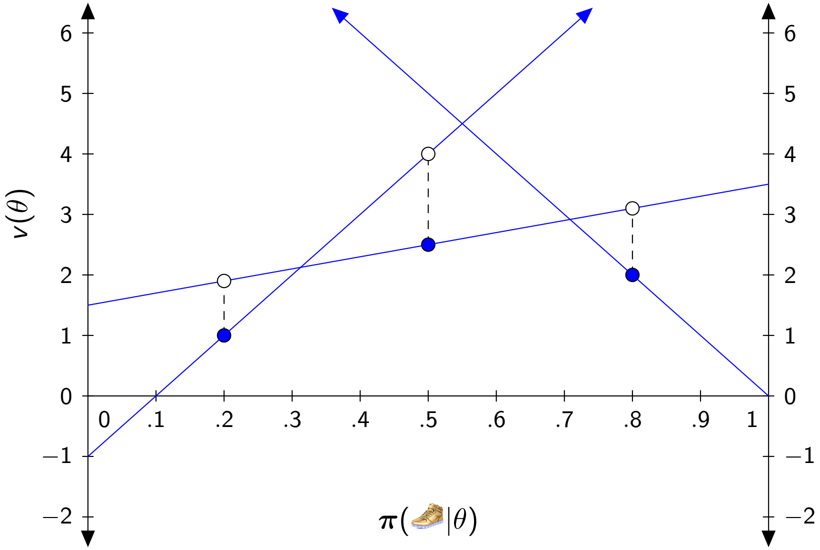

Necessary and Sufficient Condition

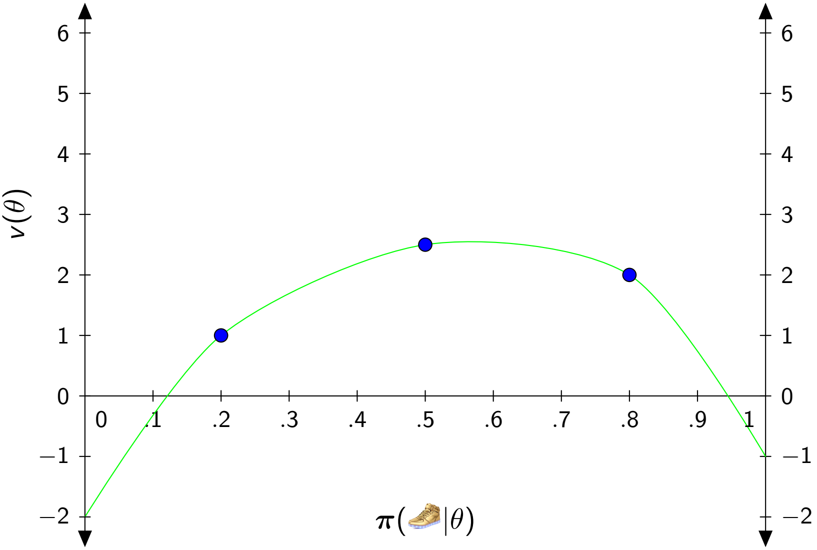

Necessary and Sufficient Condition

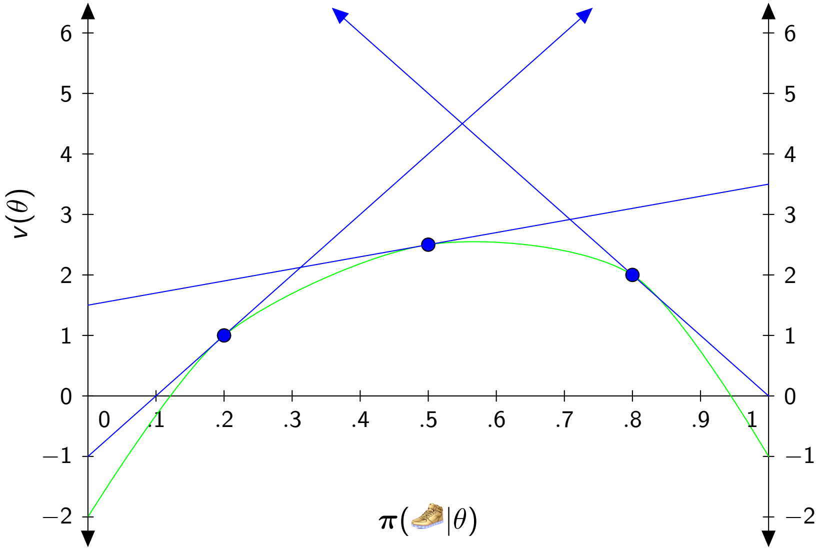

Necessary and Sufficient Condition

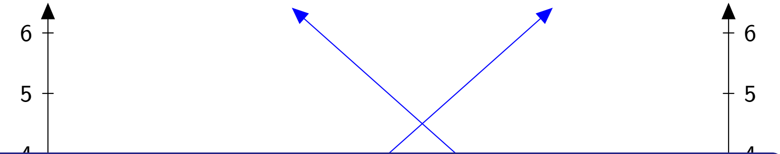

Theorem: Full Surplus Extraction with a Bayesian Mechanism (AAAI 2016)

For a given (π, Θ, Ω), full surplus extraction is possible for a Bayesian mechanism if and only if there exists a concave function $G : \mathbb{R}^{|\Omega|} \rightarrow \mathbb{R}$ such that $G(\pi(\bullet|\theta)) = v(\theta)$. ▸ Full Discussion

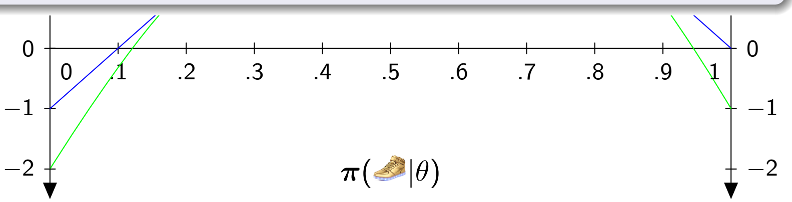

What if we have access to samples and there is "sufficient" correlation?

What if we have access to samples and there is "sufficient" correlation?

Theorem: Learning is Impossible (AAMAS 17)

For any finite number of samples, there exists a distribution for which the optimal learned mechanism is no better than the ex-post mechanism.

Discussion

Consistent Distributions

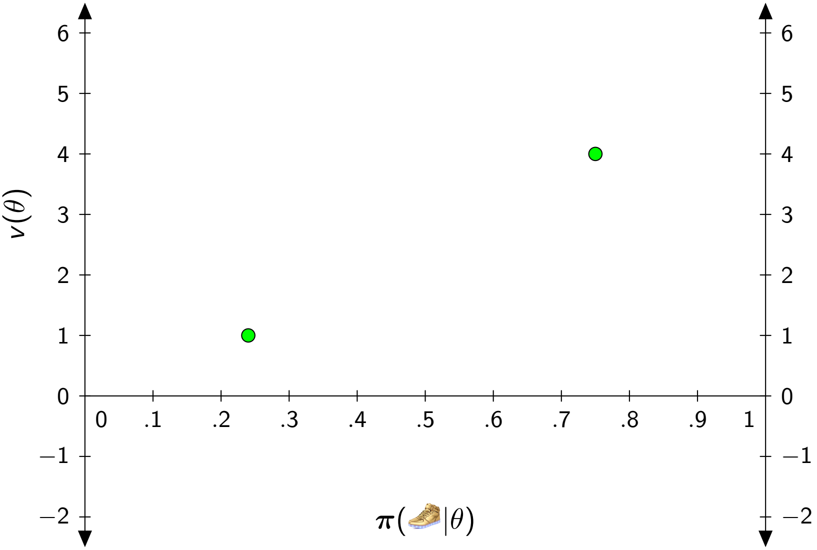

Consistent Distributions

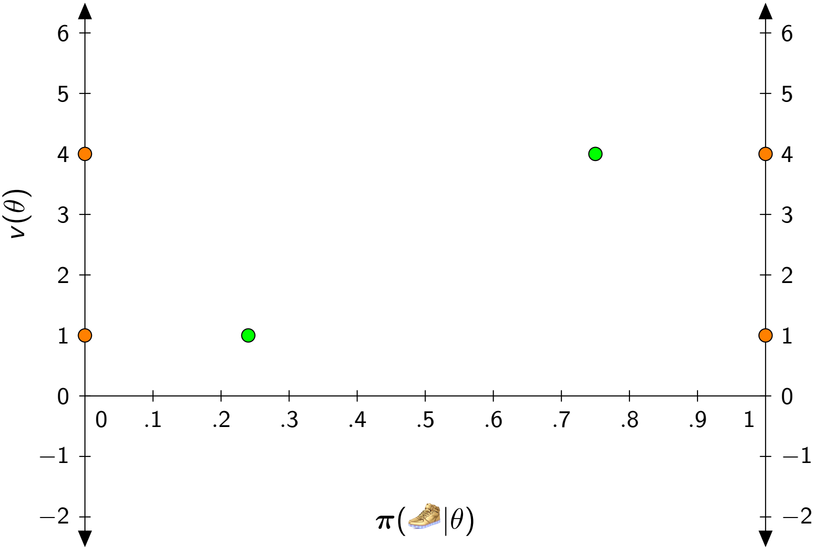

Consistent Distributions

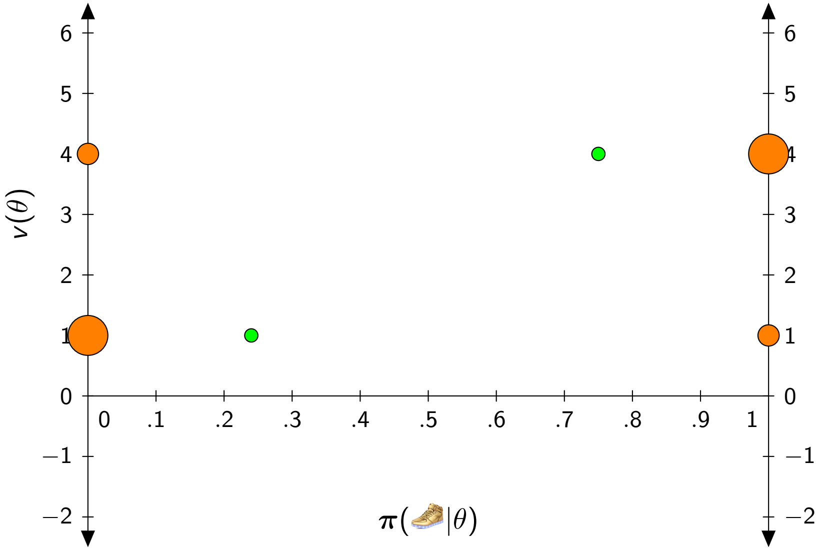

Consistent Distributions

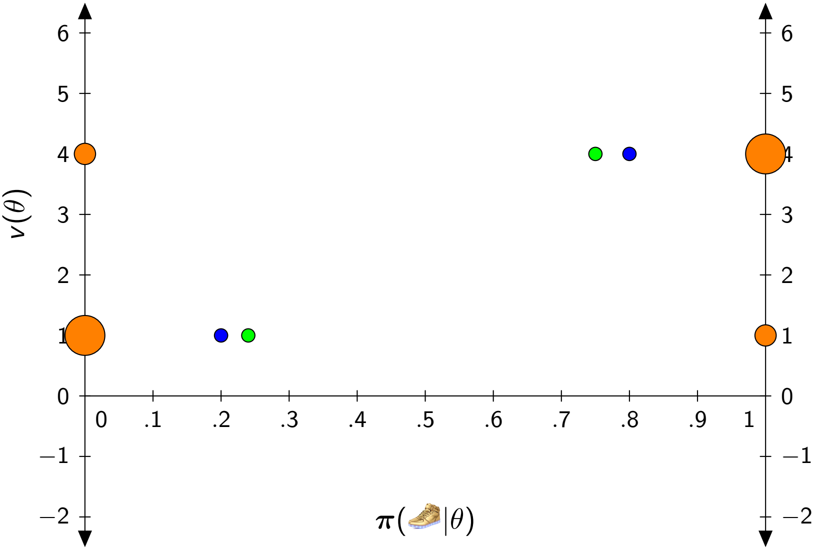

Consistent Distributions

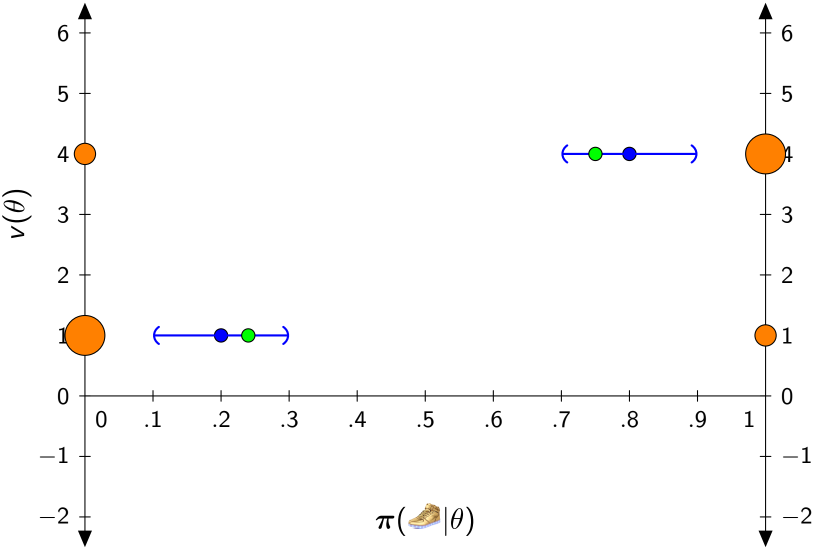

Consistent Distributions

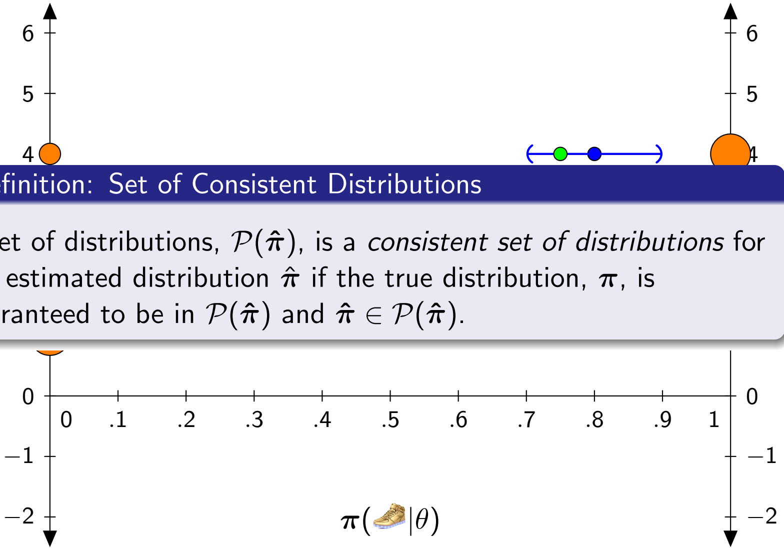

Definition: Set of Consistent Distributions

A set of distributions, $\mathcal{P}(\hat{\pi})$, 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 Mechanism Design

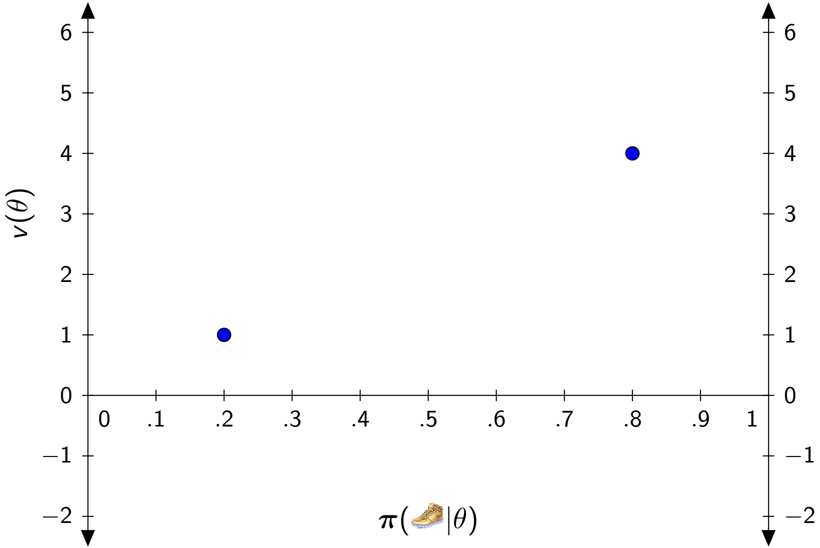

Robust Mechanism Design

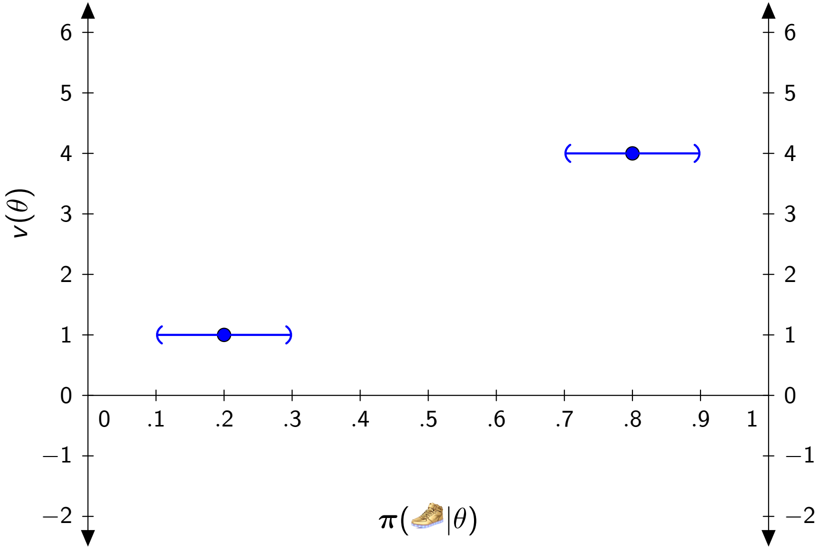

Robust Mechanism Design

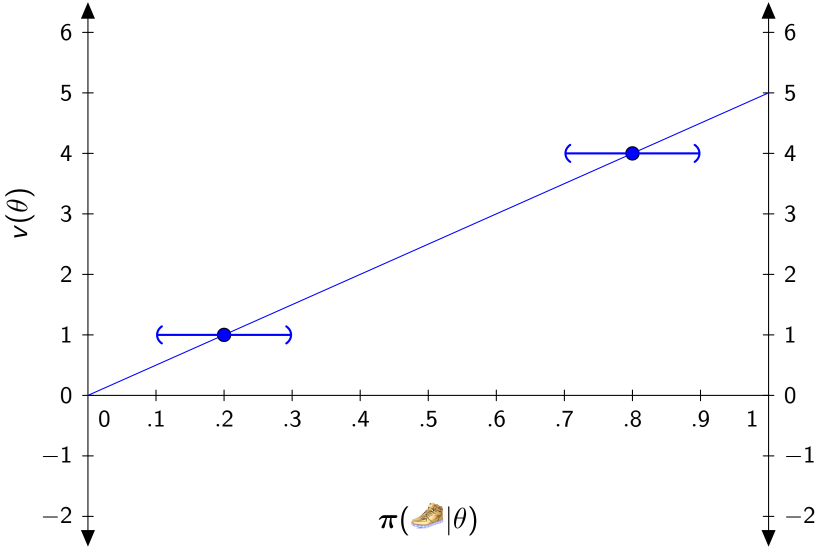

Robust Mechanism Design

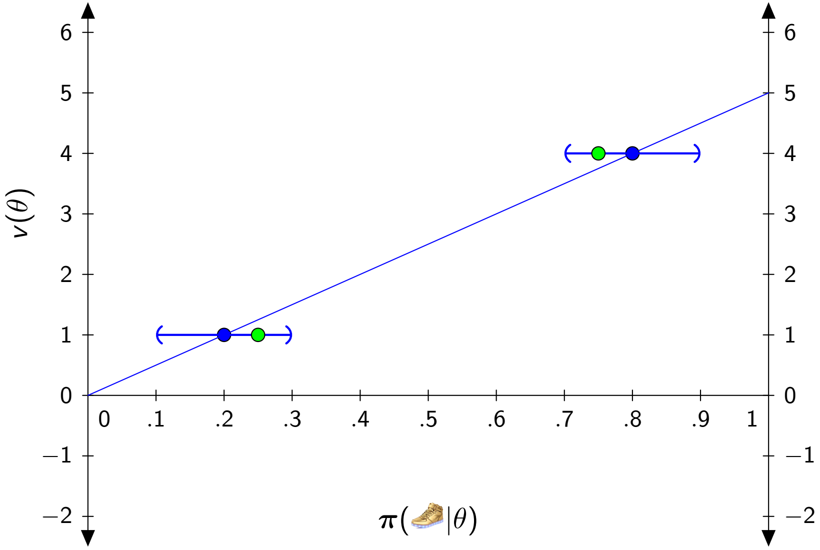

Robust Mechanism Design

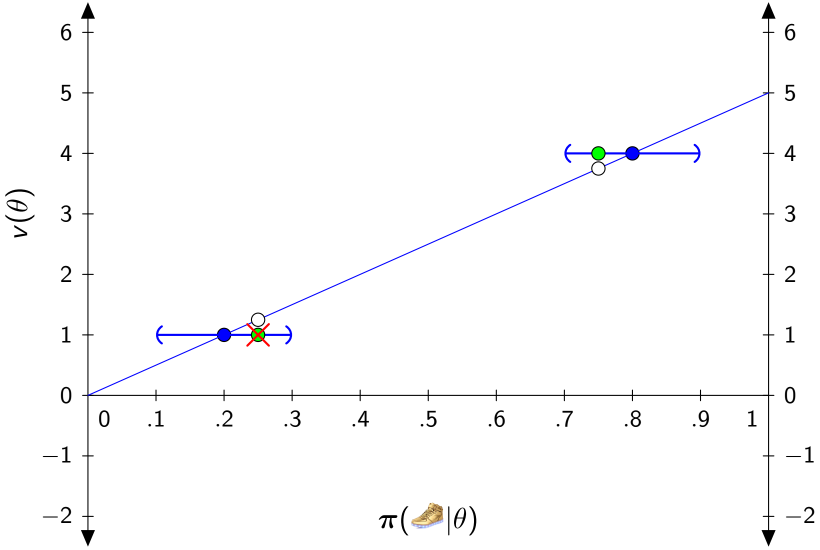

Robust Mechanism Design

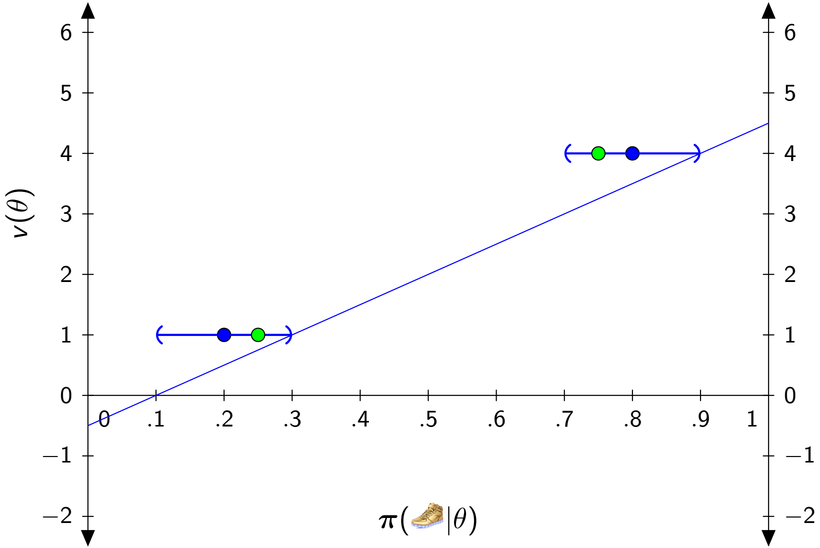

Robust Mechanism Design

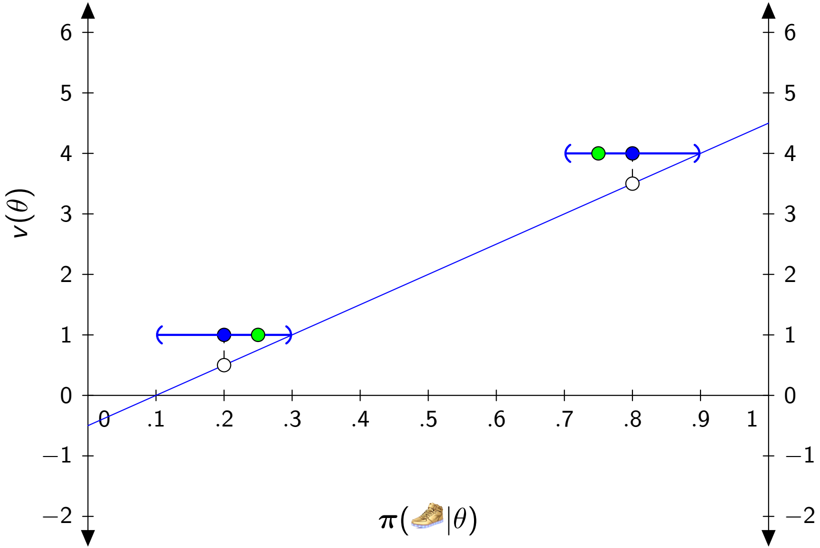

Robust Mechanism Design

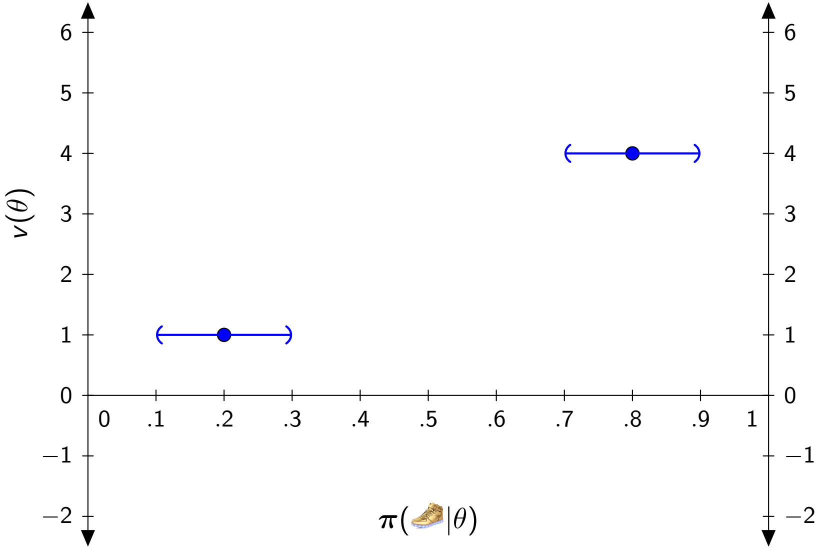

Robust Mechanism Design

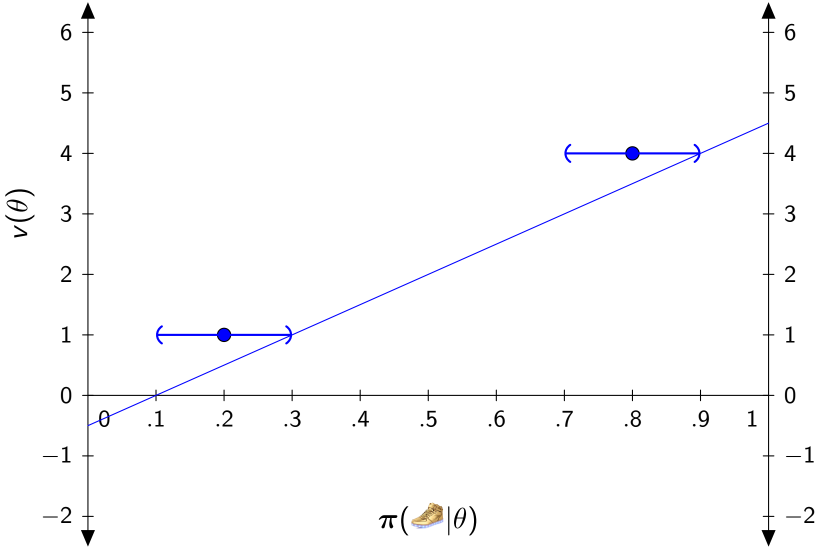

Robust Mechanism Design

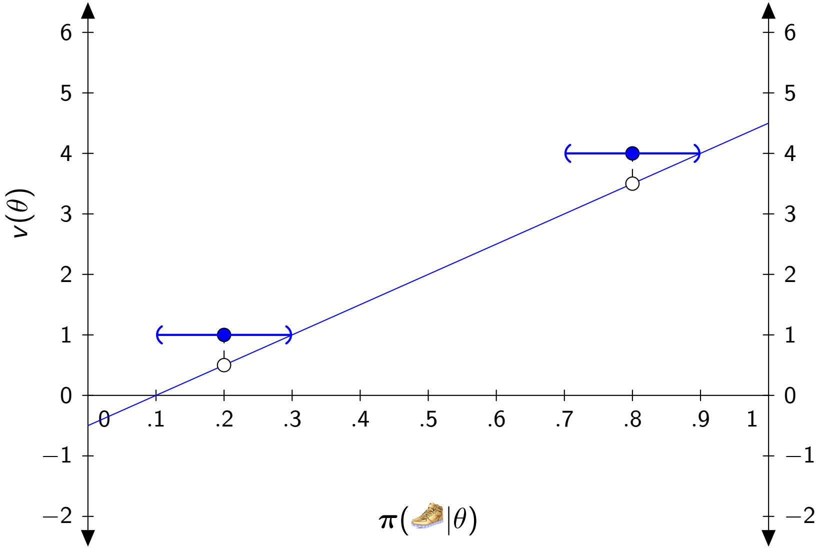

Robust Mechanism Design

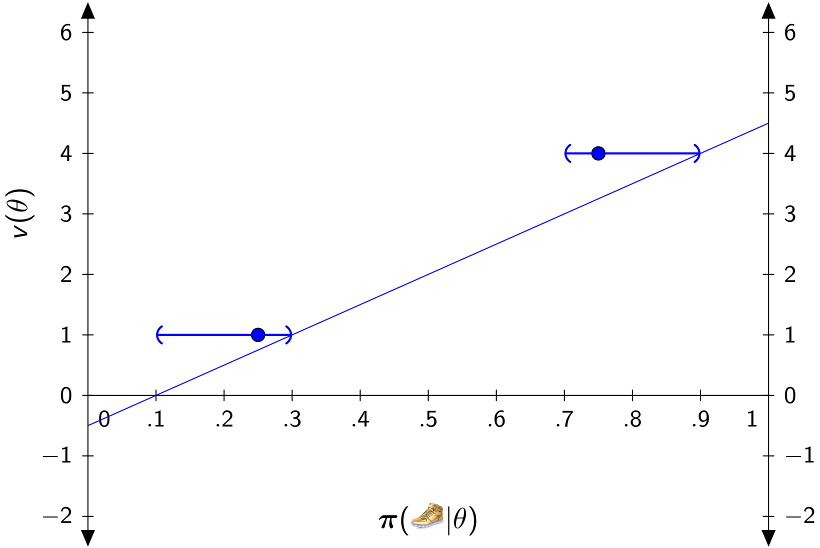

Robust Mechanism Design

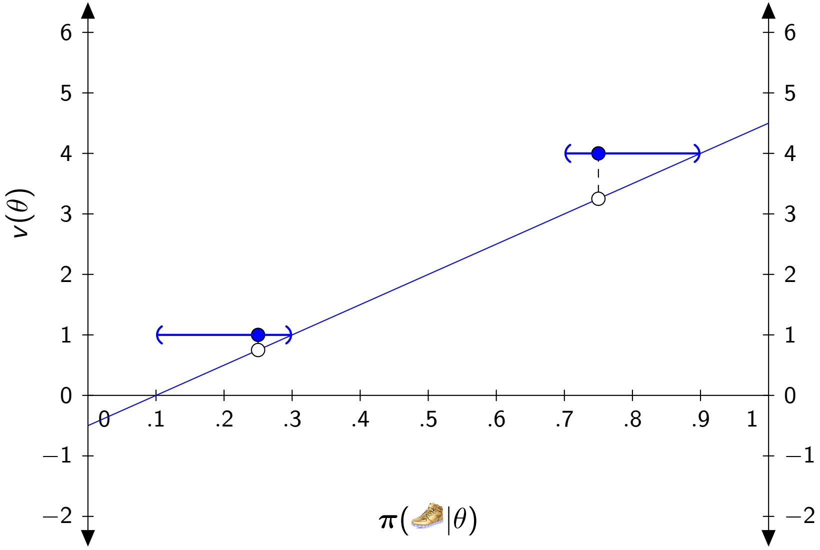

Robust Mechanism Design

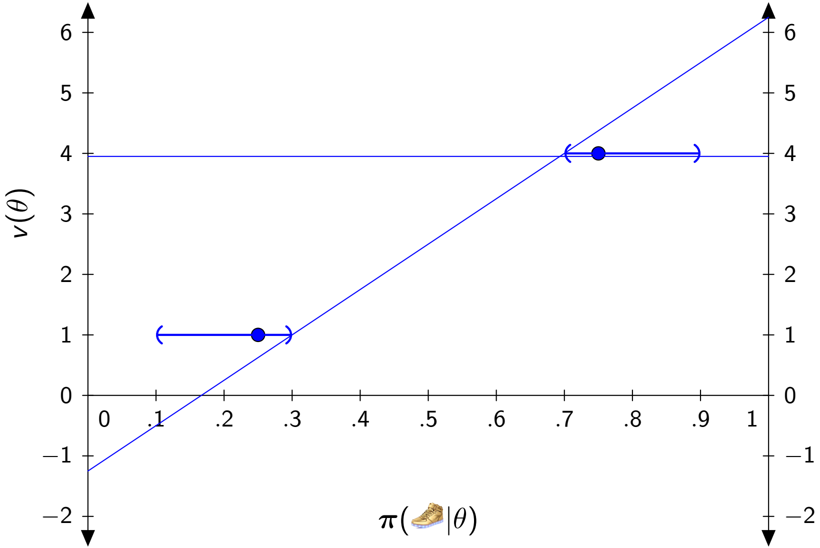

Robust Mechanism Design

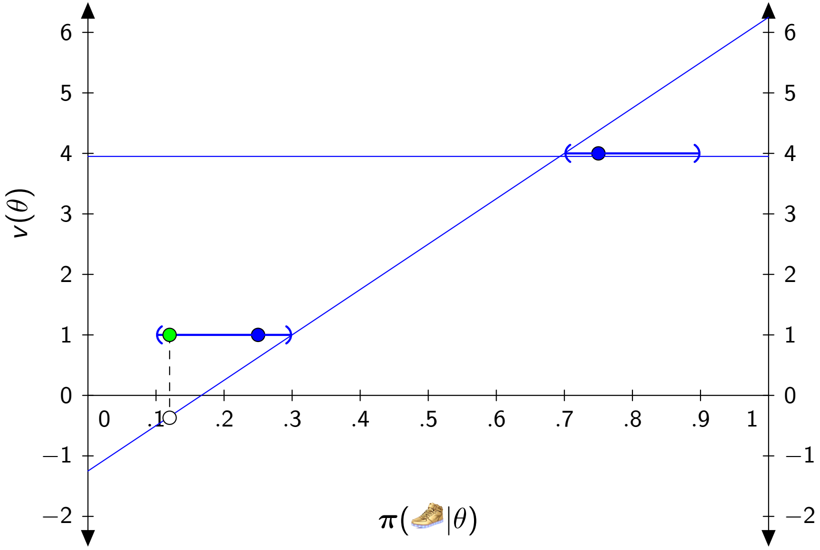

Illustration of the Algorithm

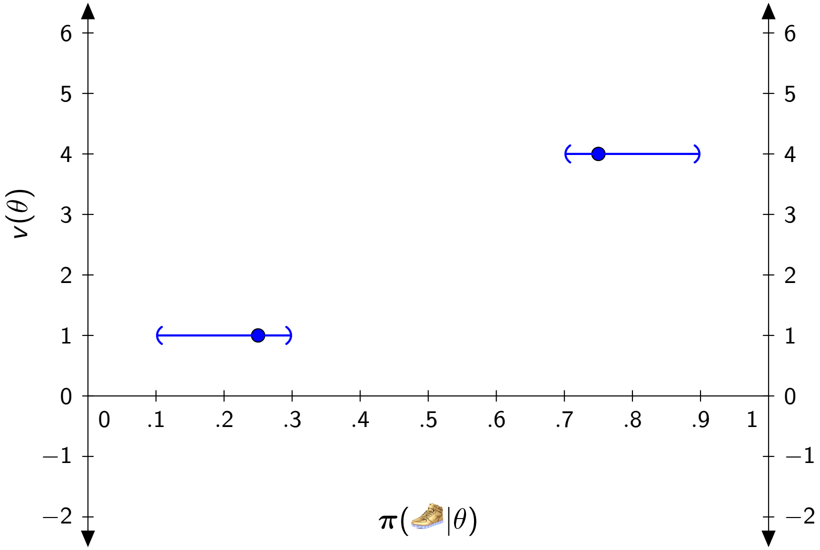

Illustration of the Algorithm

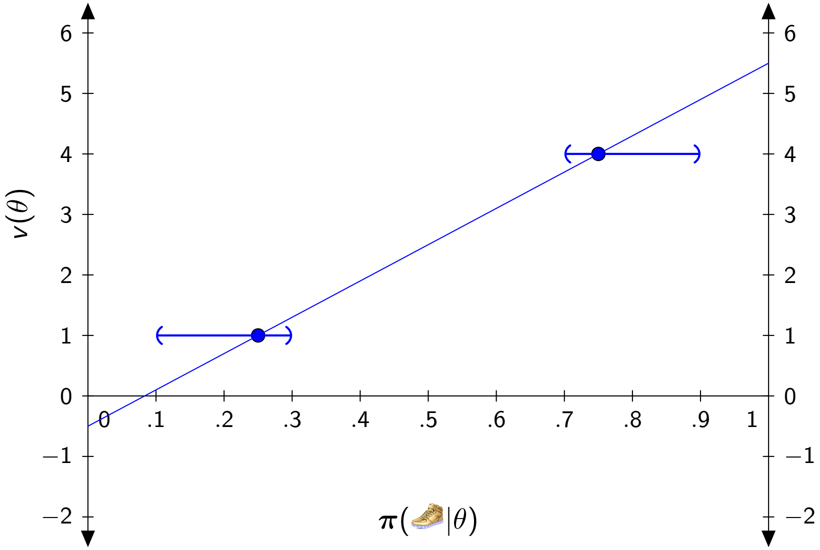

Illustration of the Algorithm

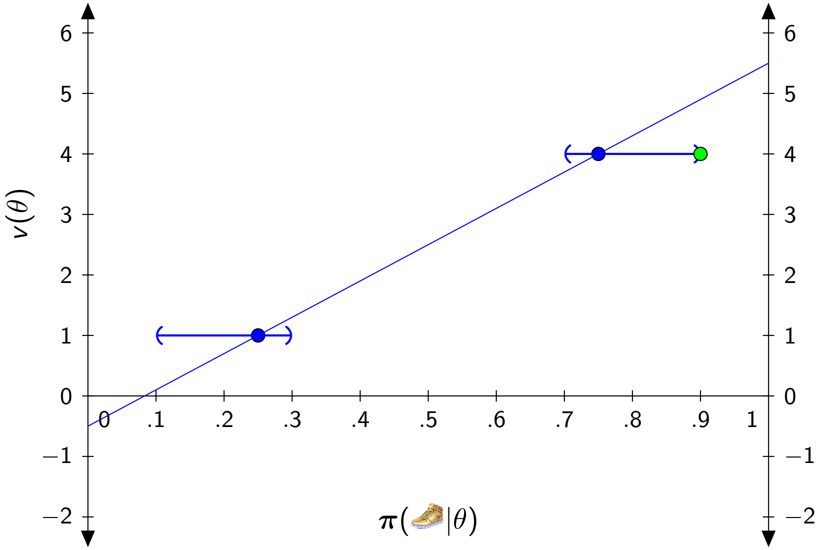

Illustration of the Algorithm

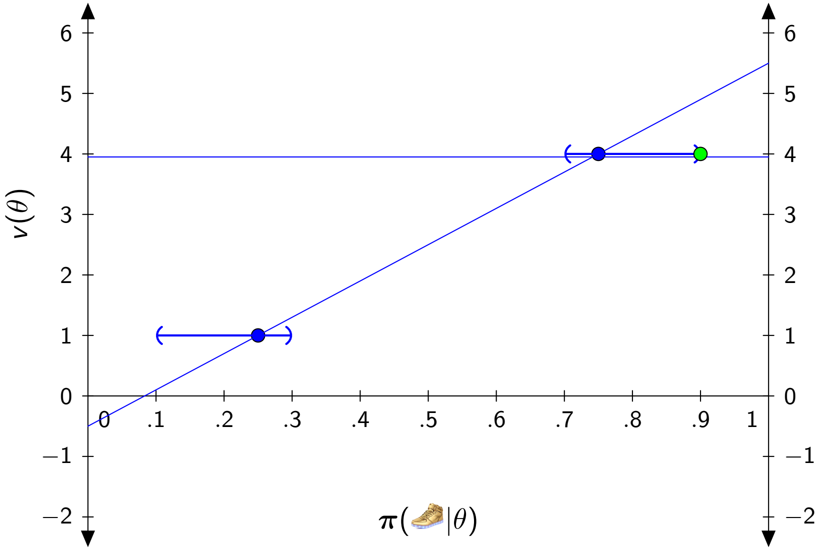

Illustration of the Algorithm

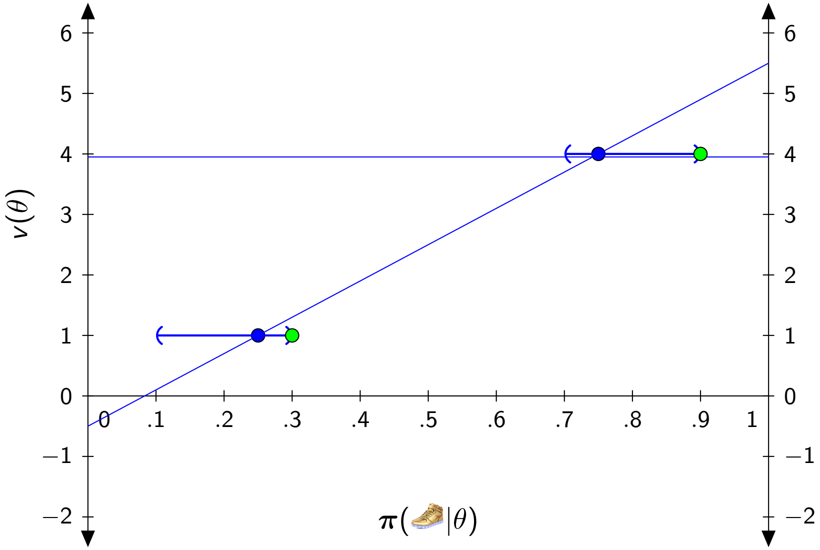

Illustration of the Algorithm

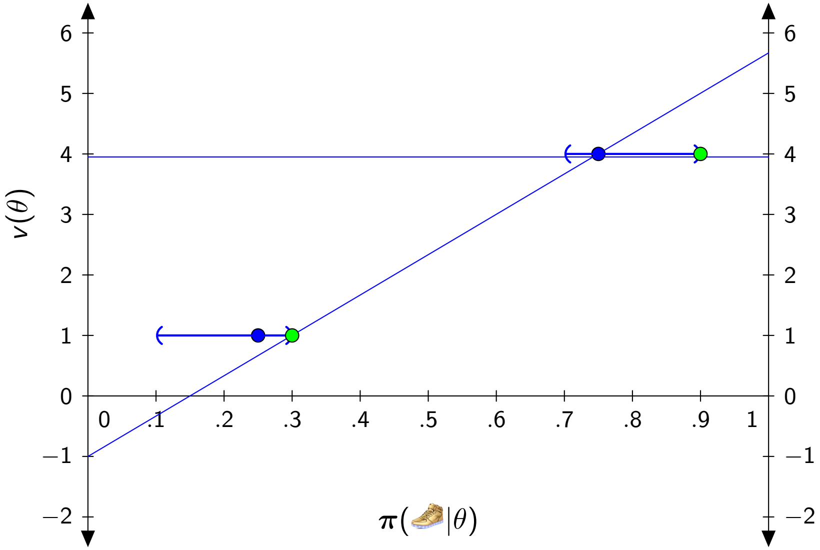

Illustration of the Algorithm

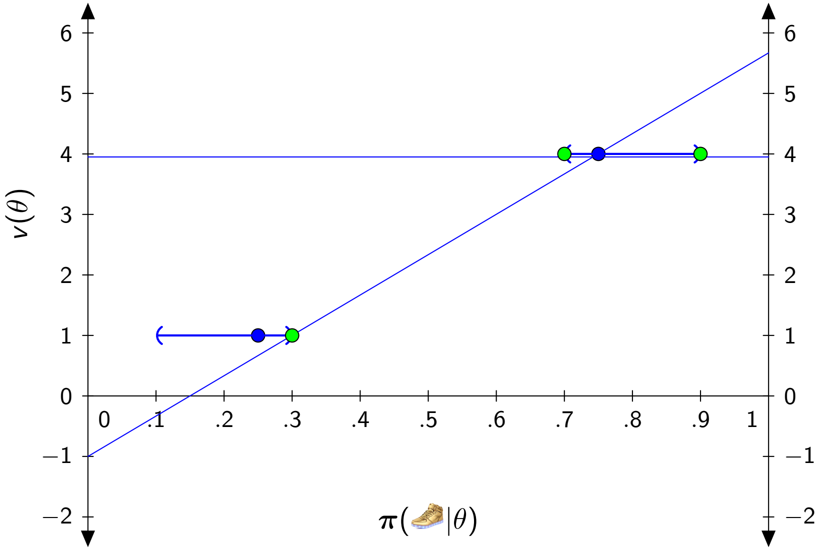

Illustration of the Algorithm

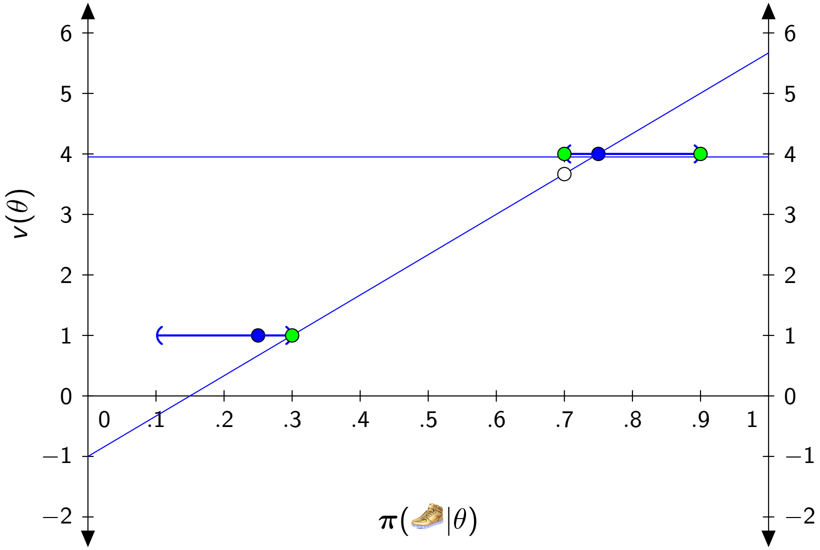

Illustration of the Algorithm

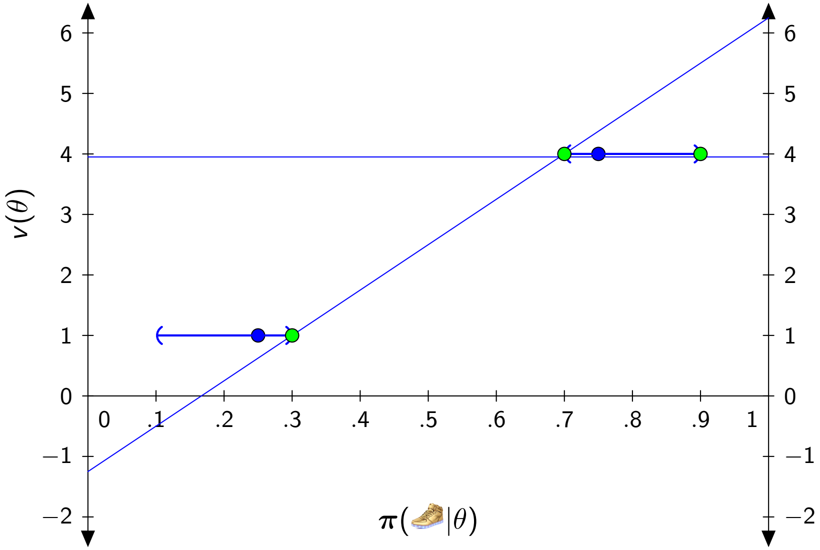

Illustration of the Algorithm

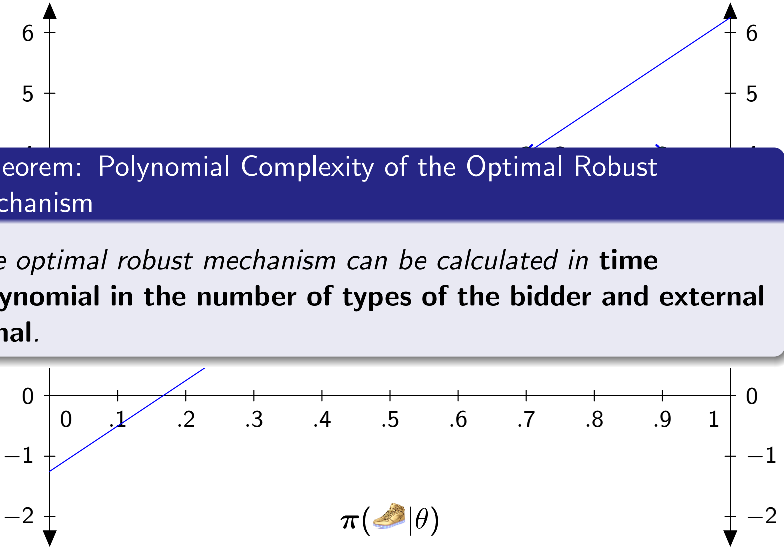

Theorem: Polynomial Complexity of the Optimal Robust Mechanism

The optimal robust mechanism can be calculated in **time polynomial in the number of types of the bidder and external signal***.*

ϵ-Robust Mechanism Design

Robust is not sufficient

- All results and intuition for 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 set of distributions, $\mathcal{P}_\epsilon(\hat{\boldsymbol{\pi}})$, is an *$\epsilon$-consistent set of distributions* for the estimated distribution $\hat{\boldsymbol{\pi}}$ if the true distribution, $\boldsymbol{\pi}$, is in $\mathcal{P}_\epsilon(\hat{\boldsymbol{\pi}})$ with probability $1-\epsilon$ and $\hat{\boldsymbol{\pi}} \in \mathcal{P}_\epsilon(\hat{\boldsymbol{\pi}})$.

- All results and intuition for 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 set of distributions, $\mathcal{P}_\epsilon(\hat{\boldsymbol{\pi}})$, is an ϵ-*consistent set of distributions* for the estimated distribution $\hat{\boldsymbol{\pi}}$ if the true distribution, $\boldsymbol{\pi}$, is in $\mathcal{P}_\epsilon(\hat{\boldsymbol{\pi}})$ with probability $1-\epsilon$ and $\hat{\boldsymbol{\pi}} \in \mathcal{P}_\epsilon(\hat{\boldsymbol{\pi}})$.

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

Parameterized Bayesian IC and IR with ϵ

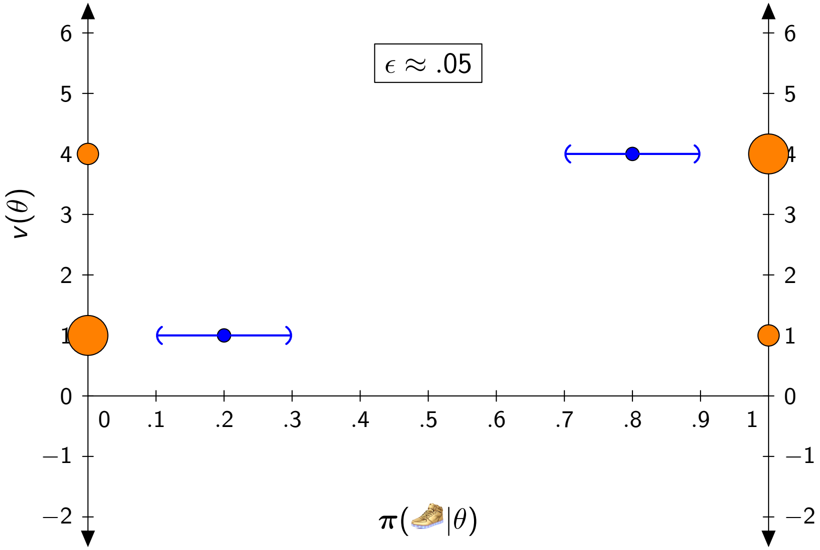

Parameterized Bayesian IC and IR with ϵ

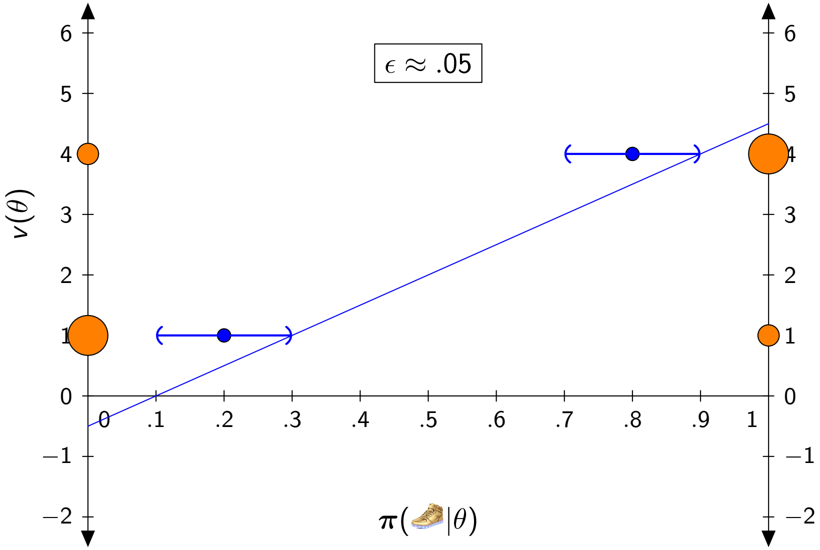

Parameterized Bayesian IC and IR with ϵ

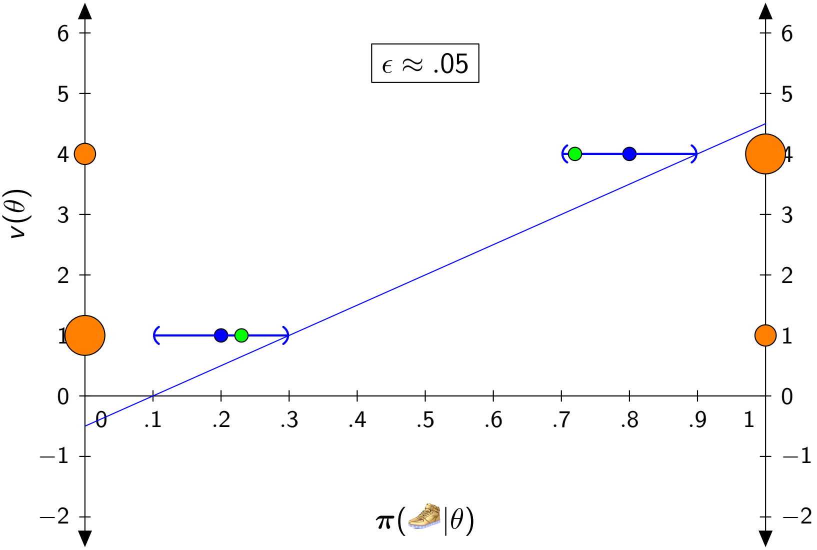

Parameterized Bayesian IC and IR with ϵ

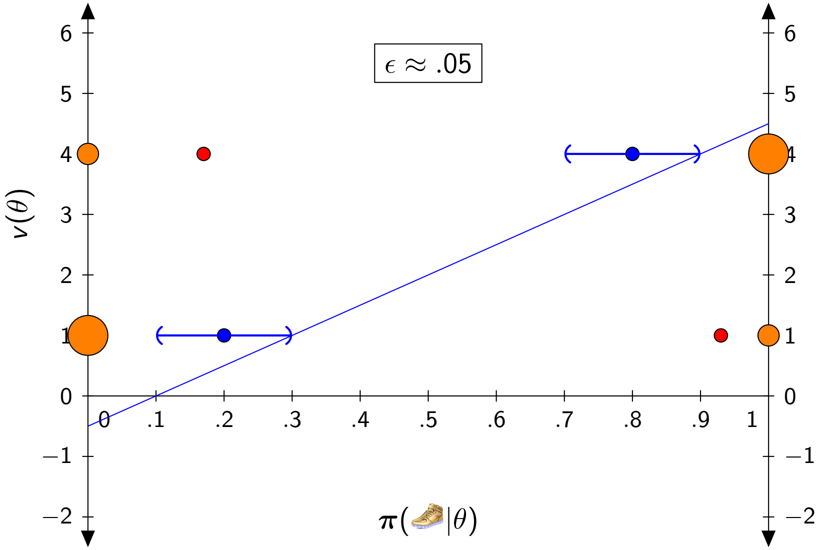

Parameterized Bayesian IC and IR with ϵ

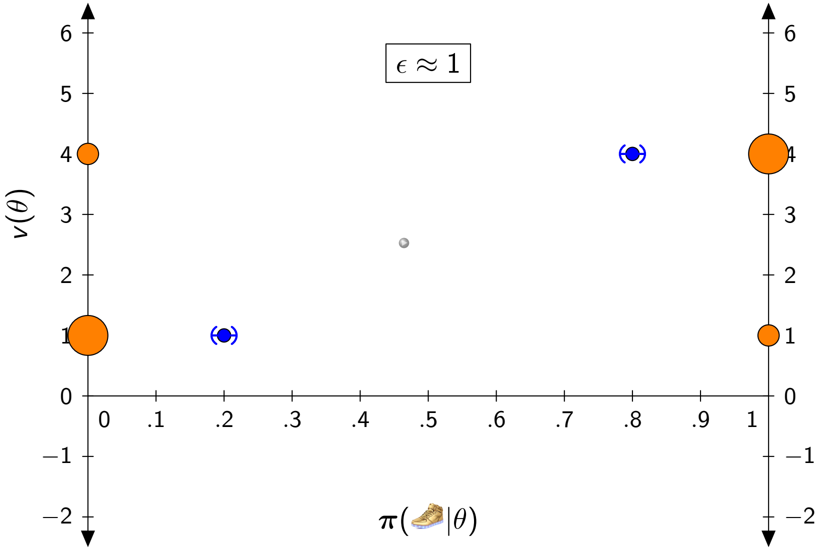

Parameterized Bayesian IC and IR with ϵ

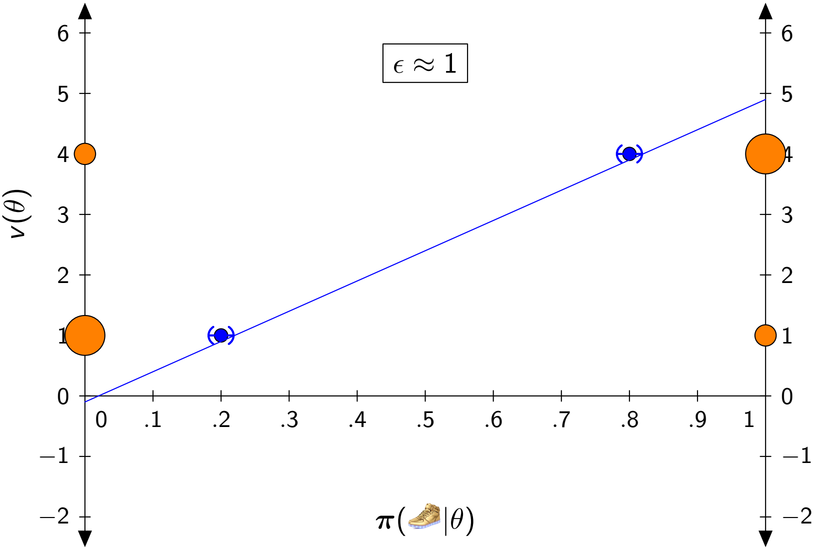

Parameterized Bayesian IC and IR with ϵ

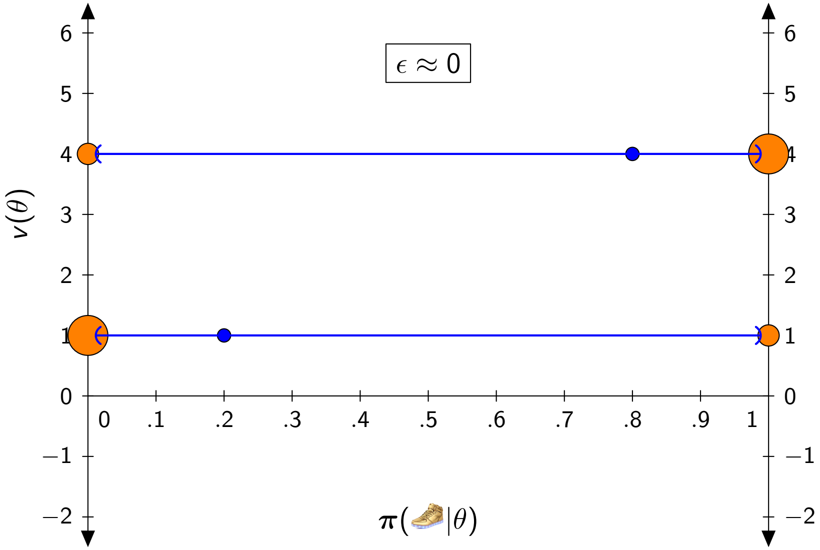

Parameterized Bayesian IC and IR with ϵ

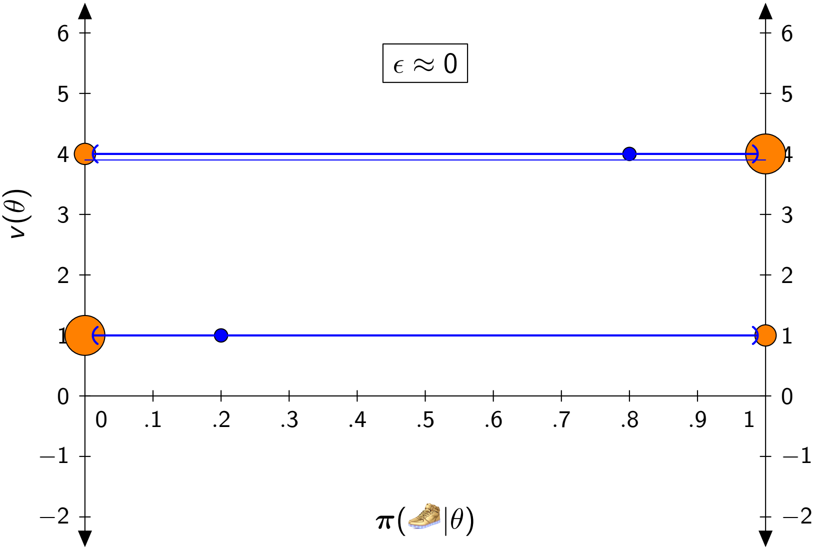

Revenue Guarantee for ϵ-Robust Mechanism Design

How do ϵ-robust mechanisms perform?

- By our impossibility result, may perform arbitrarily badly
- We require that beliefs be γ-*separated*

Revenue Guarantee for ϵ-Robust Mechanism Design

How do ϵ-robust mechanisms perform?

- By our impossibility result, may perform arbitrarily badly
- We require that beliefs be γ-*separated*

Revenue Guarantee for ϵ-Robust Mechanism Design

How do ϵ-robust mechanisms perform?

- By our impossibility result, may perform arbitrarily badly
- We require that beliefs be γ-*separated*

Revenue Guarantee for ϵ-Robust Mechanism Design

How do ϵ-robust mechanisms perform?

- By our impossibility result, may perform arbitrarily badly
- We require that beliefs be γ-*separated*

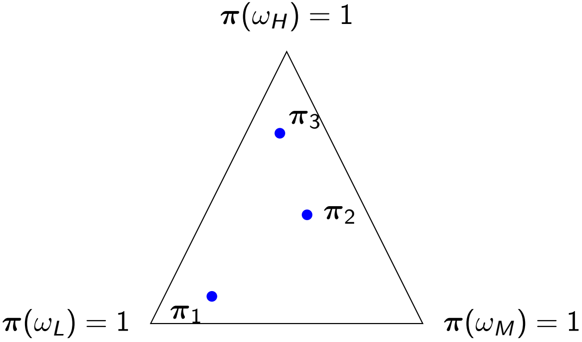

Revenue Guarantee for ϵ-Robust Mechanism Design

How do ϵ-robust mechanisms perform?

- By our impossibility result, may perform arbitrarily badly
- We require that beliefs be γ-*separated*

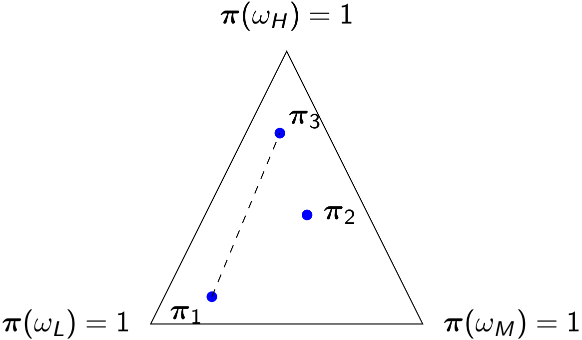

Revenue Guarantee for ϵ-Robust Mechanism Design

How do ϵ-robust mechanisms perform?

- By our impossibility result, may perform arbitrarily badly
- We require that beliefs be γ-*separated*

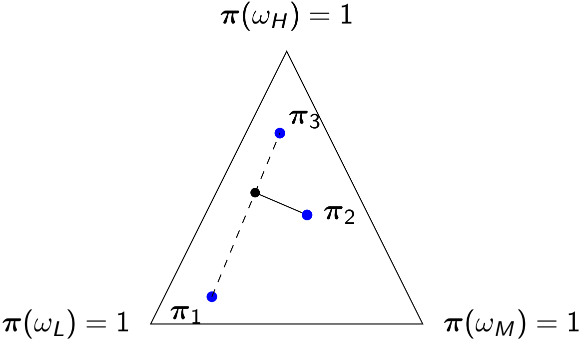

Revenue Guarantee for ϵ-Robust Mechanism Design

How do ϵ-robust mechanisms perform?

- By our impossibility result, may perform arbitrarily badly
- We require that beliefs be γ-*separated*

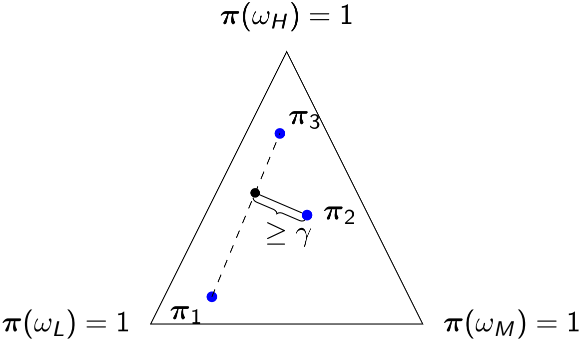

Revenue Guarantee for ϵ-Robust Mechanism Design

How do ϵ-robust mechanisms perform?

- By our impossibility result, may perform arbitrarily badly
- We require that beliefs be γ-*separated*

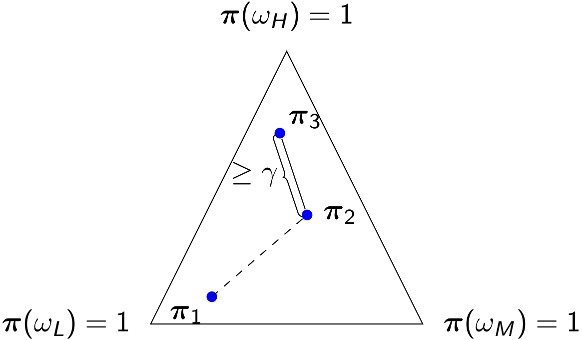

Revenue Guarantee for ϵ-Robust Mechanism Design

How do ϵ-robust mechanisms perform?

- By our impossibility result, may perform arbitrarily badly
- We require that beliefs be γ-*separated*

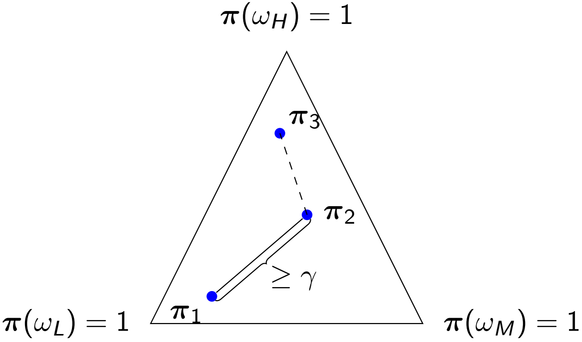

Revenue Guarantee for ϵ-Robust Mechanism Design

How do ϵ-robust mechanisms perform?

- By our impossibility result, may perform arbitrarily badly
- We require that beliefs be γ-*separated*

Sample Complexity of ϵ-Robust Mechanism Design

The **sample complexity** *for constructing an ϵ-robust mechanism that is an* **additive** *k*-**approximation** *to the optimal mechanism is* $\mathbf{O}(\mathbf{poly}(\frac{1}{\mathbf{k}}, \frac{1}{\gamma}, |\Theta|, |\Omega|, \mathbf{v}(|\Theta|)))$. *Proof*

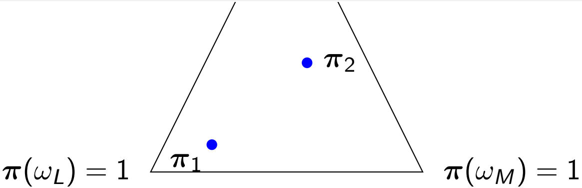

Simulations

- 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
  - $\epsilon = .05$
  - $\Theta = \{1, 2, ..., 10\}$
  - $|\Omega| = 10$
  - $v(\theta) = \theta$

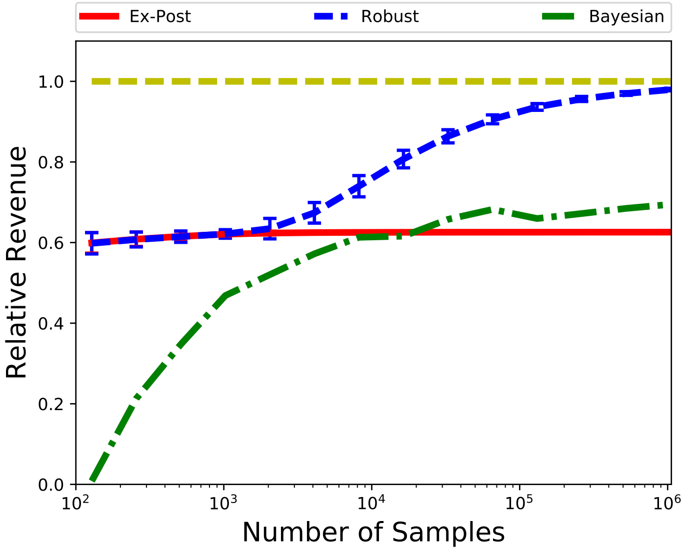
Ex-Post
Robust
Bayesian
Relative Revenue
1.0
0.8
0.6
0.4
0.2
0.0
10^2
10^3
10^4
10^5
10^6
Number of Samples

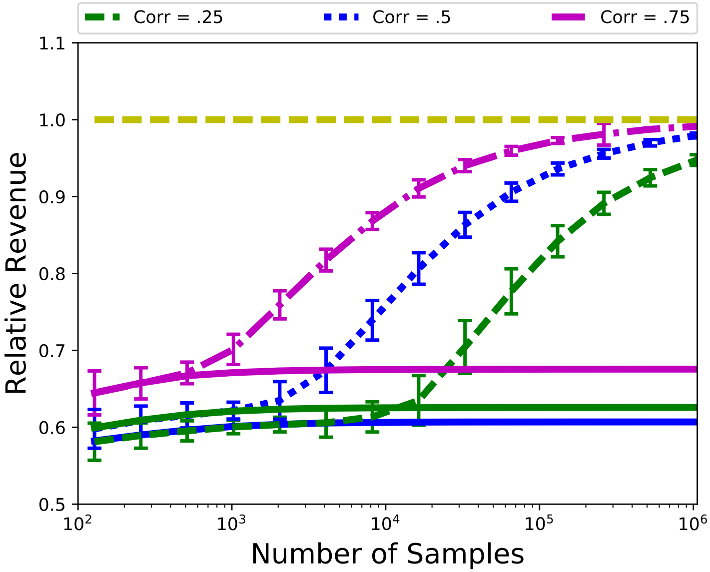
Corr = .25
Corr = .5
Corr = .75
1.1
1.0
0.9
0.8
0.7
0.6
0.5
Relative Revenue
10^2
10^3
10^4
10^5
10^6
Number of Samples

Related Work for Robust Mechanism Design

- Robust mechanisms (Bergemann and Morris 2005)
- 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)

Full Revenue Extraction with an Ex-Post IC Mechanism

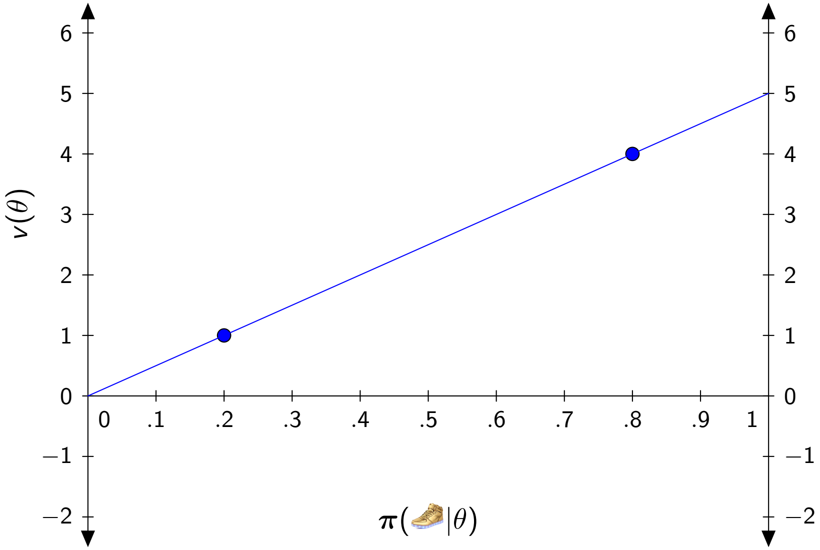

Full Revenue Extraction with an Ex-Post IC Mechanism

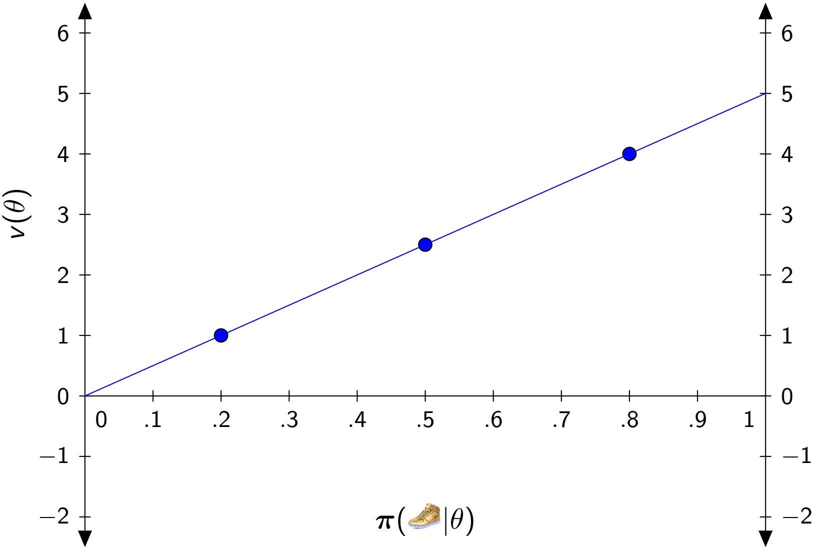

Full Revenue Extraction with an Ex-Post IC Mechanism

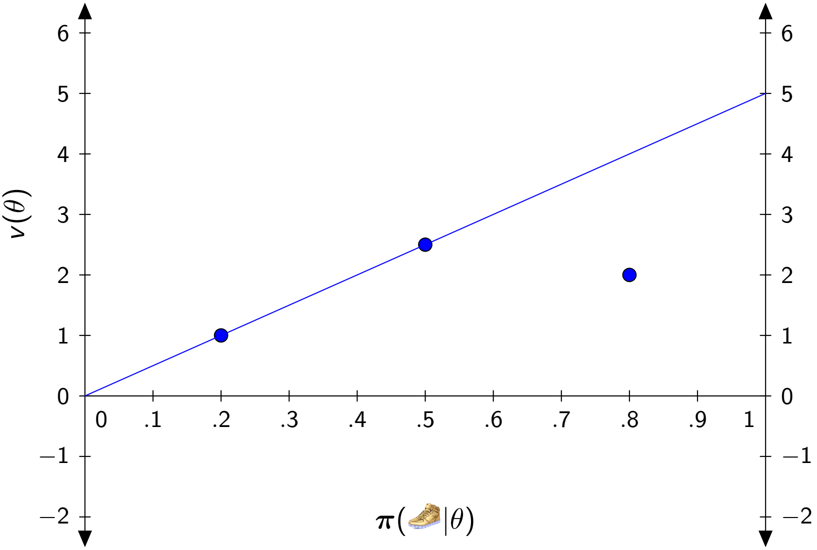

Full Revenue Extraction with an Ex-Post IC Mechanism

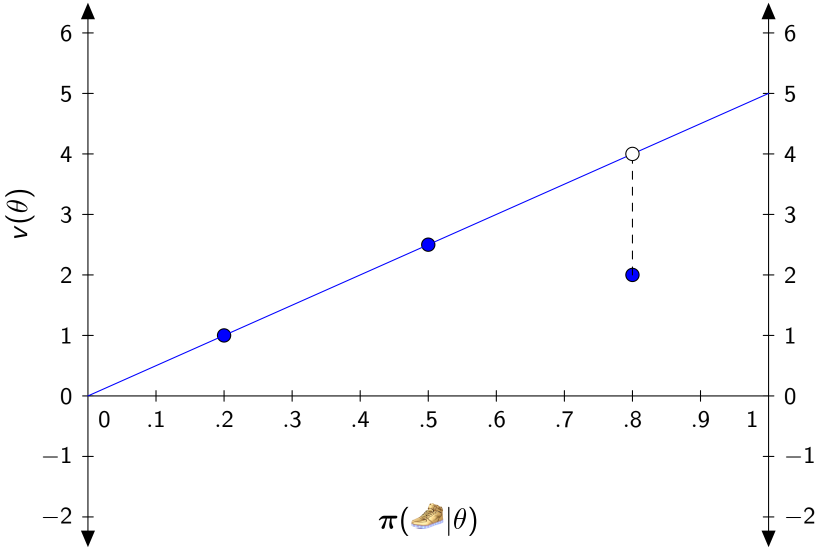

Full Revenue Extraction with an Ex-Post IC Mechanism

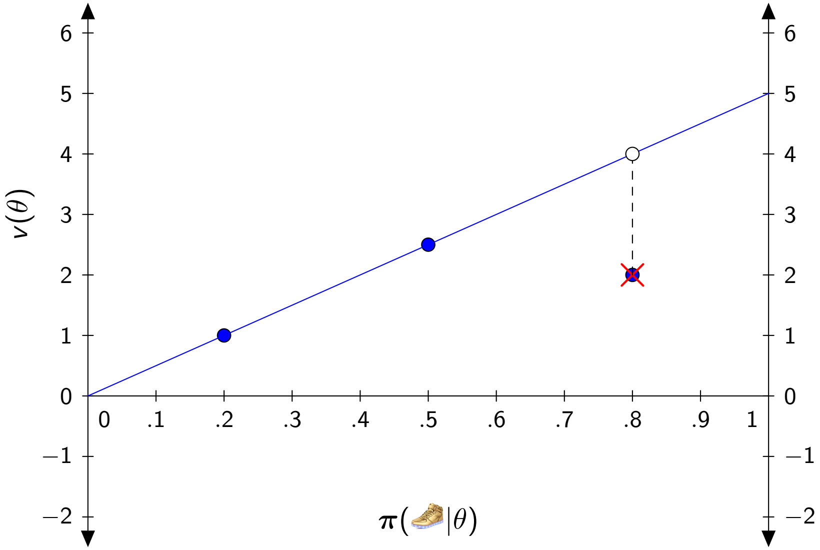

Full Revenue Extraction with an Ex-Post IC Mechanism

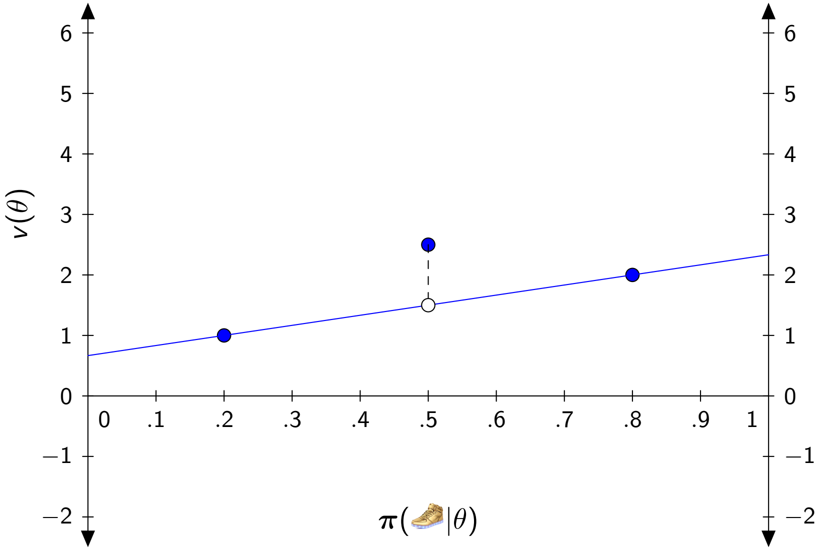

Full Revenue Extraction with an Ex-Post IC Mechanism

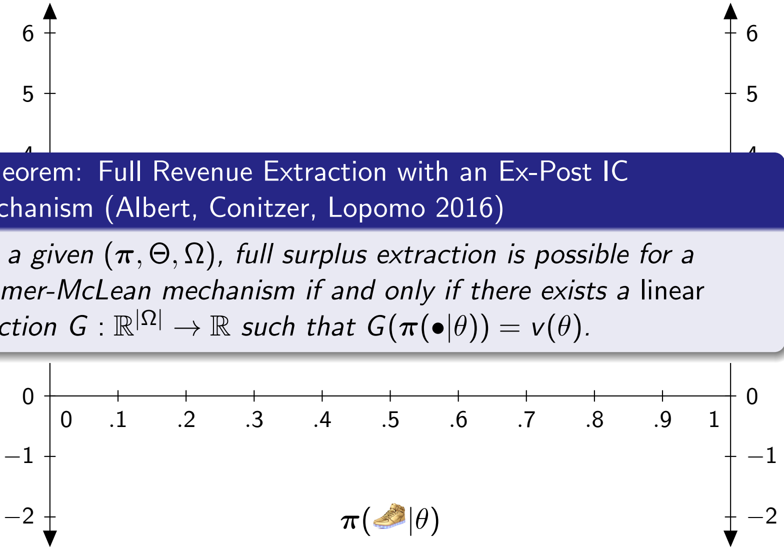

Theorem: Full Revenue Extraction with an Ex-Post IC Mechanism (Albert, Conitzer, Lopomo 2016)

For a given (π, Θ, Ω), full surplus extraction is possible for a Cremer-McLean mechanism if and only if there exists a linear *function $G : \mathbb{R}^{|\Omega|} \to \mathbb{R}$ such that $G(\pi(\bullet|\theta)) = v(\theta)$.*

Full Revenue Extraction with a Bayesian Mechanism

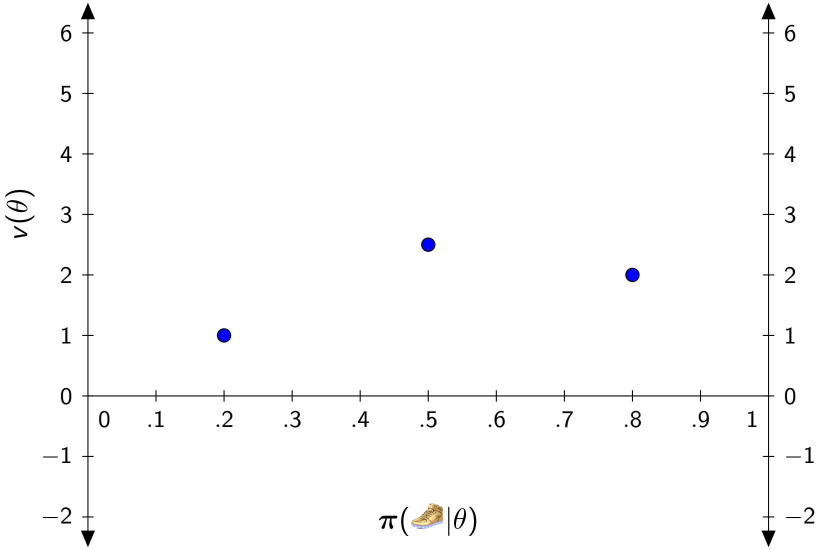

Full Revenue Extraction with a Bayesian Mechanism

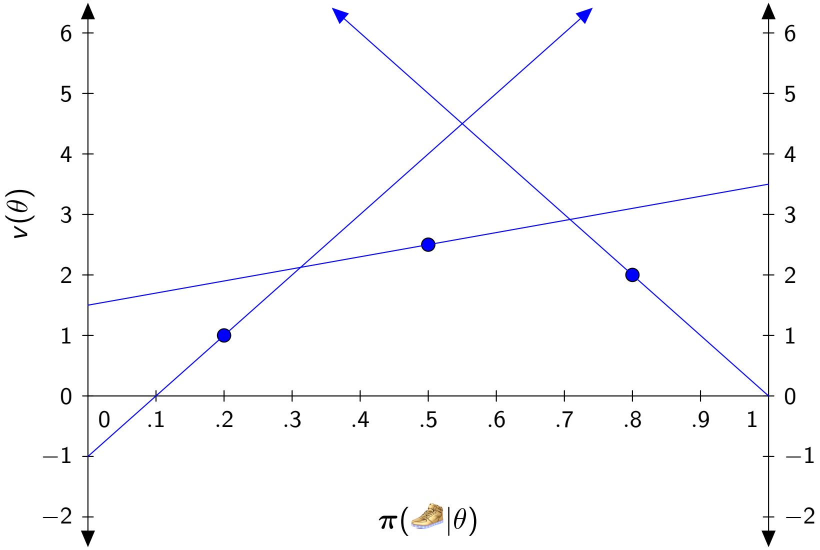

Full Revenue Extraction with a Bayesian Mechanism

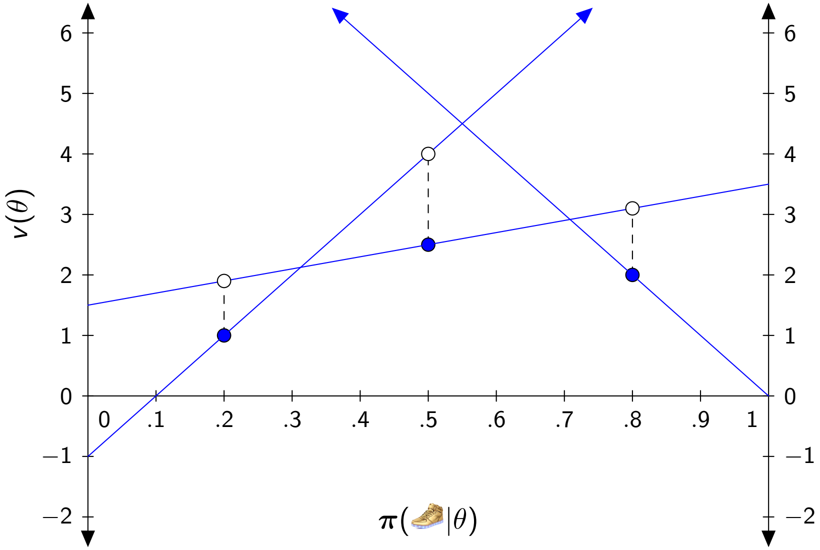

Full Revenue Extraction with a Bayesian Mechanism

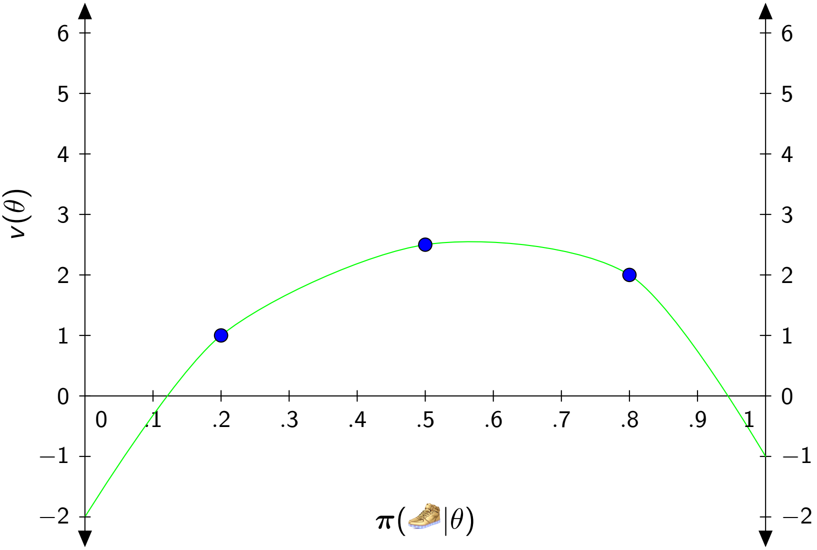

Full Revenue Extraction with a Bayesian Mechanism

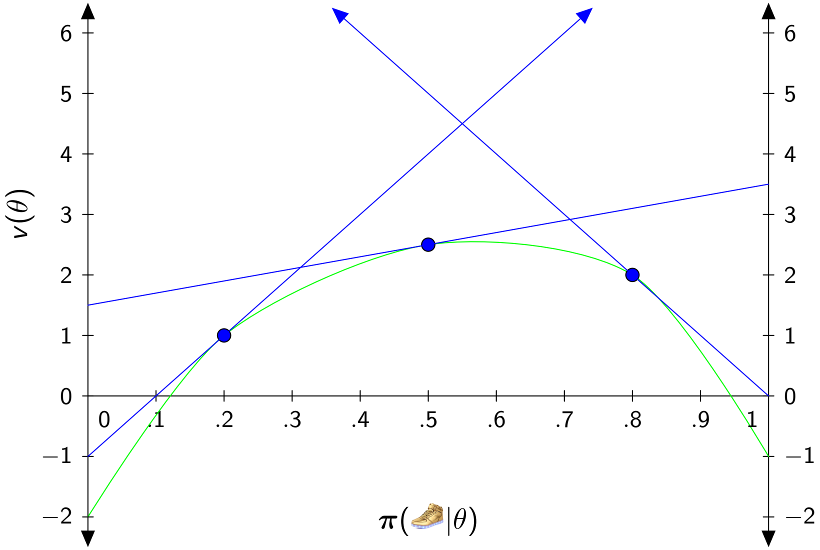

Full Revenue Extraction with a Bayesian Mechanism

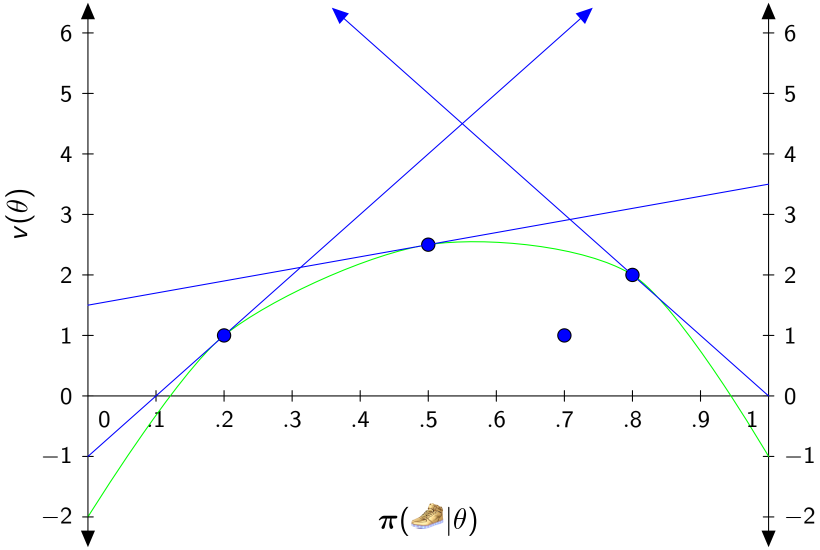

Full Revenue Extraction with a Bayesian Mechanism

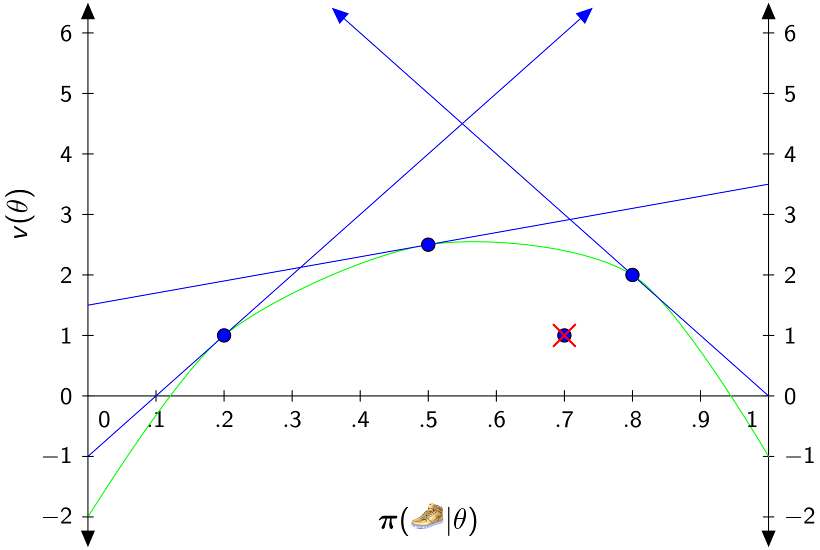

Full Revenue Extraction with a Bayesian Mechanism

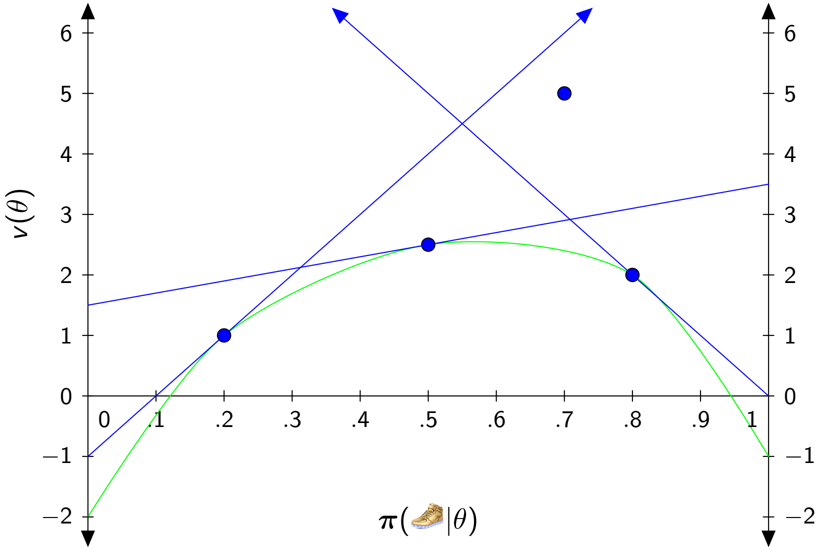

Full Revenue Extraction with a Bayesian Mechanism

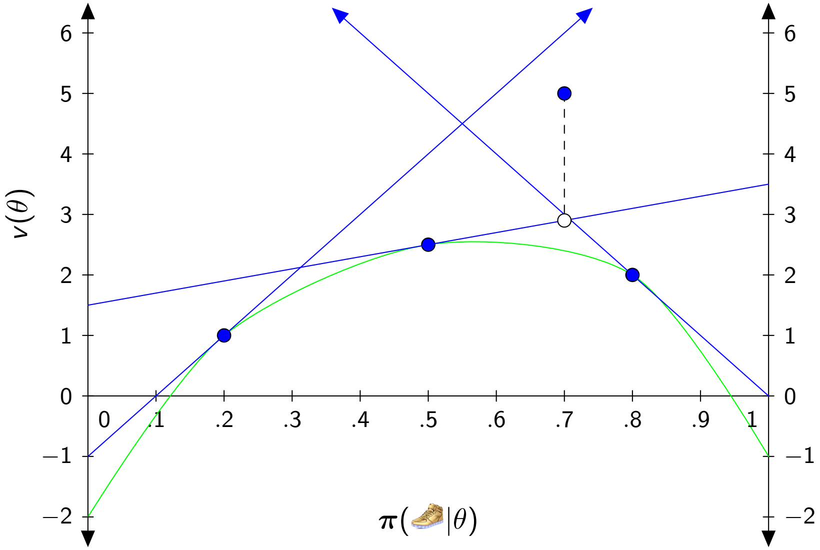

Full Revenue Extraction with a Bayesian Mechanism

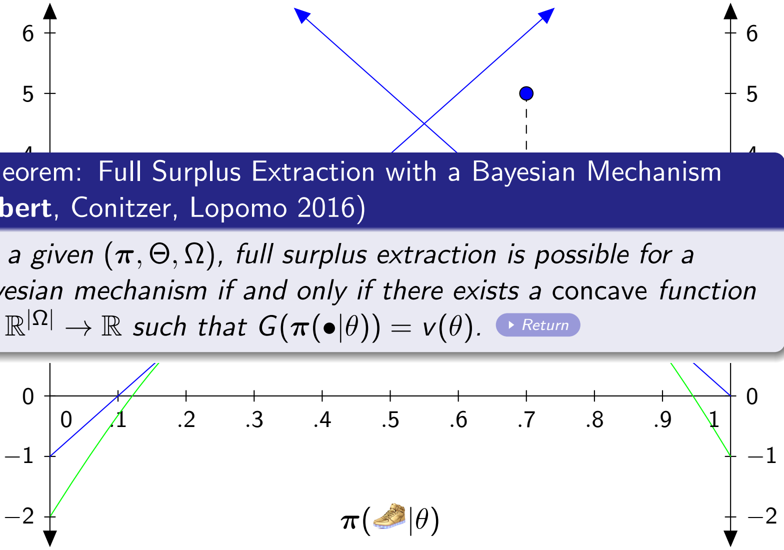

Theorem: Full Surplus Extraction with a Bayesian Mechanism (Albert, Conitzer, Lopomo 2016)

For a given (π, Θ, Ω), full surplus extraction is possible for a Bayesian mechanism if and only if there exists a concave function $G : \mathbb{R}^{|\Omega|} \to \mathbb{R}$ such that $G(\pi(\bullet|\theta)) = v(\theta)$.

Return

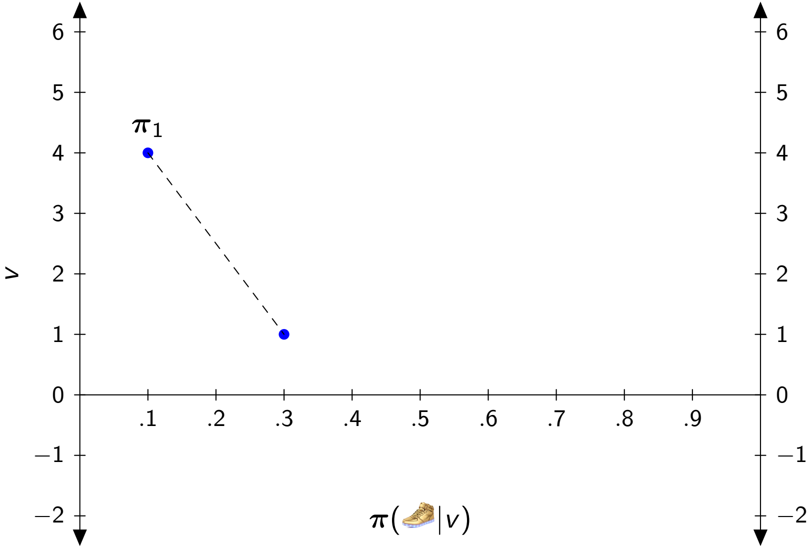
6
5
4
3
2
1
0
−1
−2
v
π_1
.1
.2
.3
.4
.5
.6
.7
.8
.9
$\pi($ $|v)$

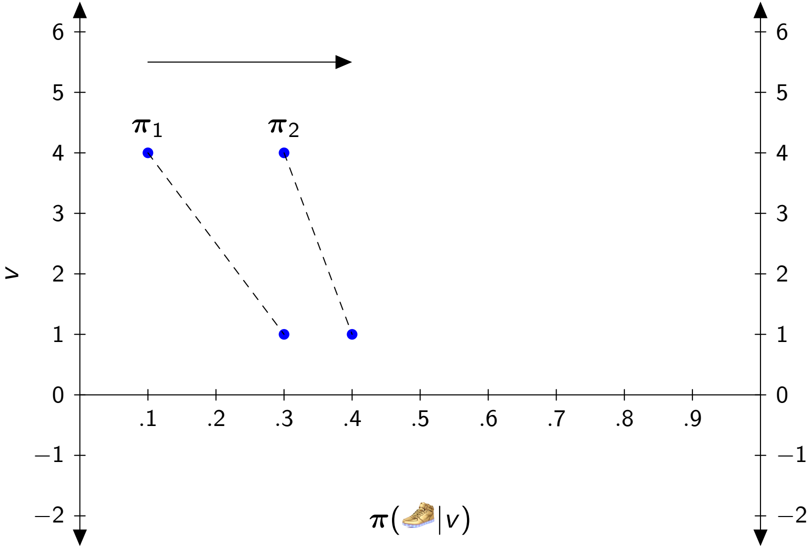
6
5
4
3
2
1
0
−1
−2
v
π_1
π_2
.1
.2
.3
.4
.5
.6
.7
.8
.9
$\pi($ | $v)$

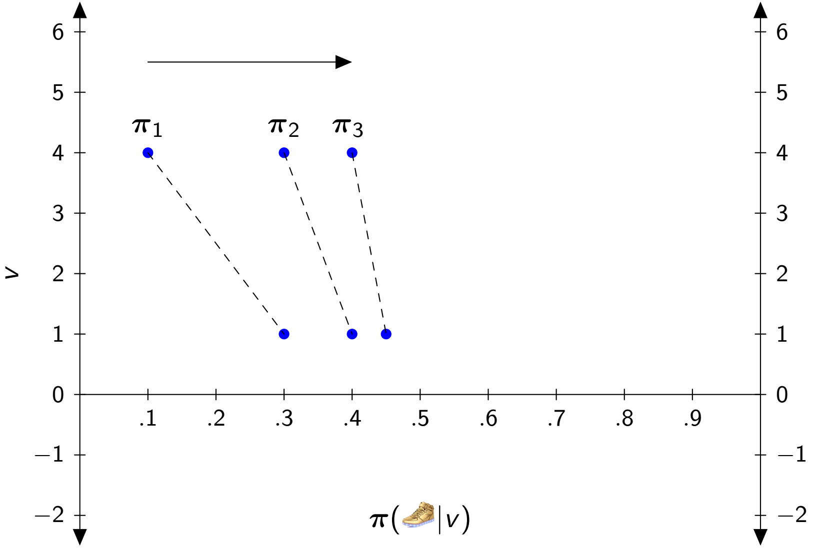

6
5
4
3
2
1
0
−1
−2
v
π1
π2
π3
.1
.2
.3
.4
.5
.6
.7
.8
.9
π(|v)

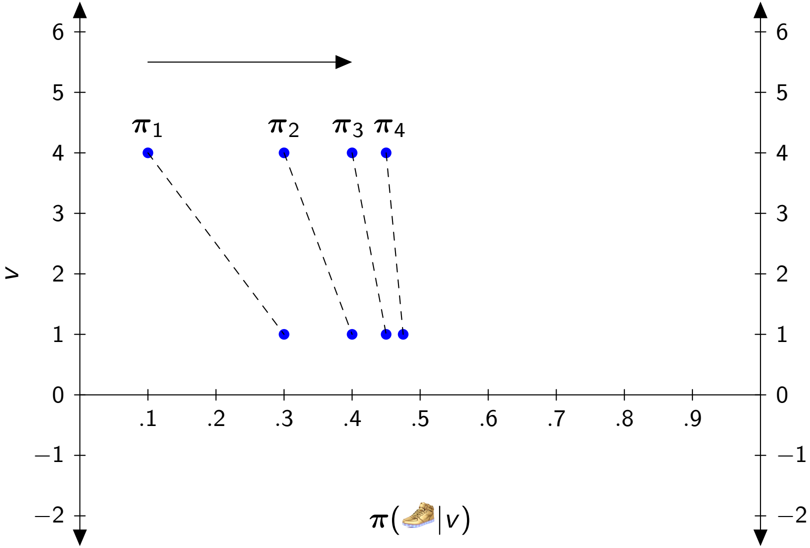
6
5
4
3
2
1
0
−1
−2
v
π_1
π_2
π_3
π_4
.1
.2
.3
.4
.5
.6
.7
.8
.9
$\pi(\cdot|v)$

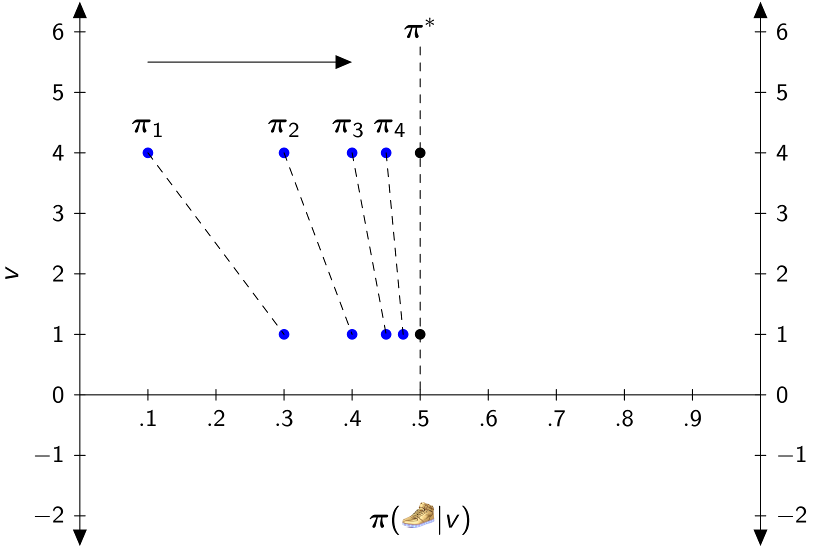
π*
π1
π2
π3
π4
6
5
4
3
2
1
0
−1
−2
v
.1
.2
.3
.4
.5
.6
.7
.8
.9
π(|v)

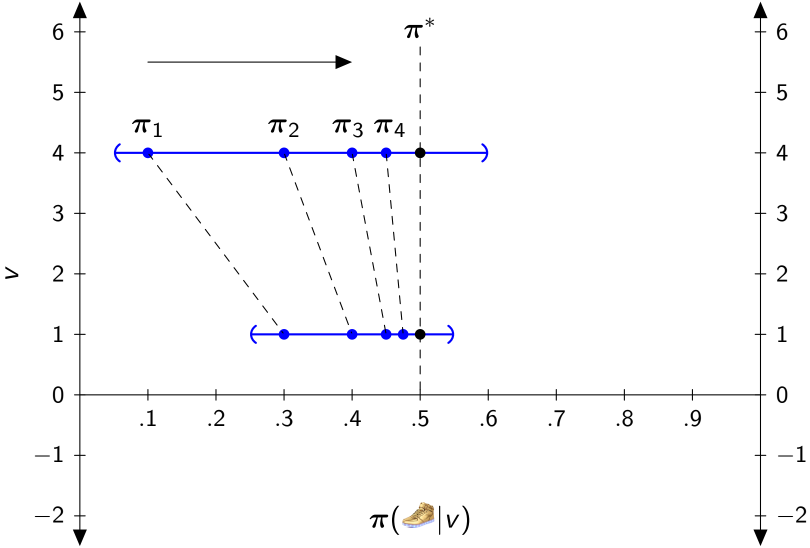

6
5
4
3
2
1
0
−1
−2
v
π*
π1
π2
π3
π4
.1
.2
.3
.4
.5
.6
.7
.8
.9
π(|v)

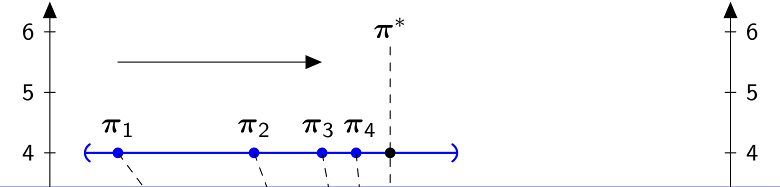

Theorem: Inapproximability of the Optimal Mechanism

For any k $>$ 0 and for any mechanism that uses a finite number of independent samples *from the underlying distribution, there exists a* $T \in \mathbb{N}$ *such that for all* $\pi_{i'} \in \{\pi_i\}_{i=T}^{\infty}$*, the expected revenue is less than* $\mathbf{OPT}(\pi^*) + \mathbf{k}$.

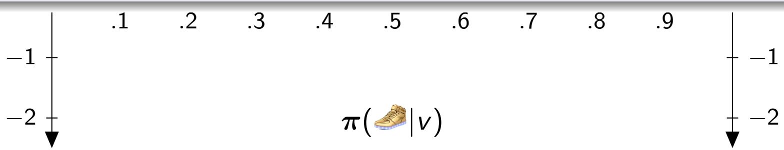

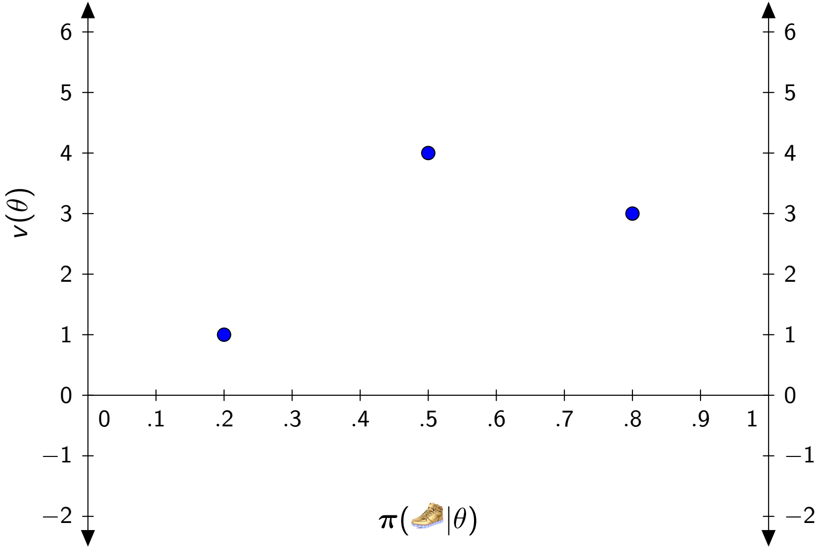

6
5
4
3
2
1
0
−1
−2
$v(\theta)$
0 .1 .2 .3 .4 .5 .6 .7 .8 .9 1
$\pi($👟$|\theta)$

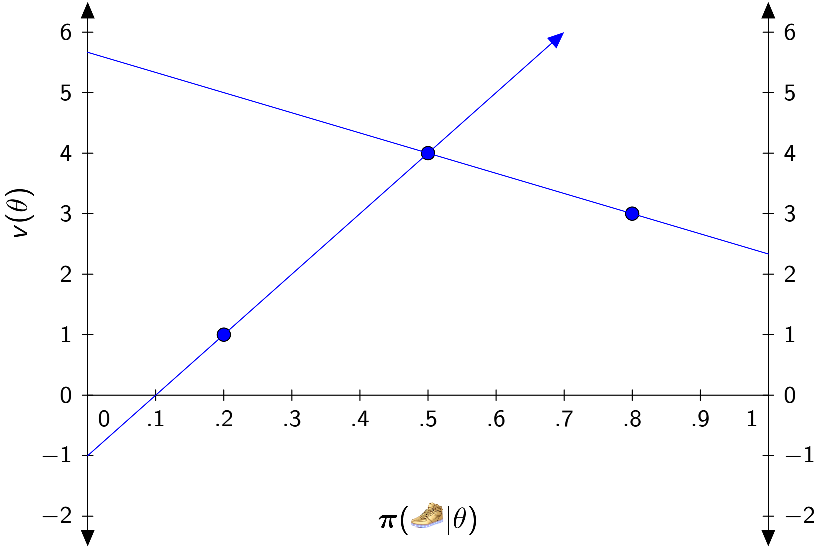
6
5
4
3
2
1
0
−1
−2
0 .1 .2 .3 .4 .5 .6 .7 .8 .9 1
$v(\theta)$
$\pi(\text{👟}|\theta)$

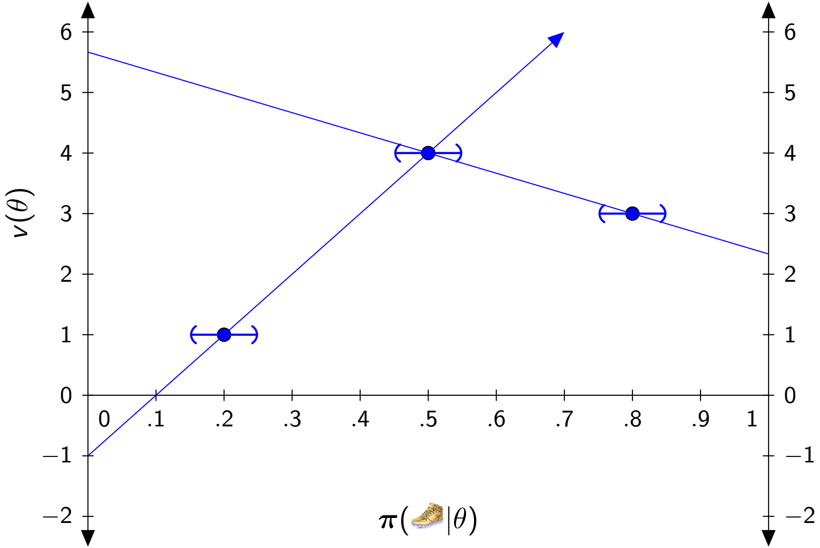
6
5
4
3
2
1
0
−1
−2
$v(\theta)$
0
.1
.2
.3
.4
.5
.6
.7
.8
.9
1
$\pi(\ |\theta)$

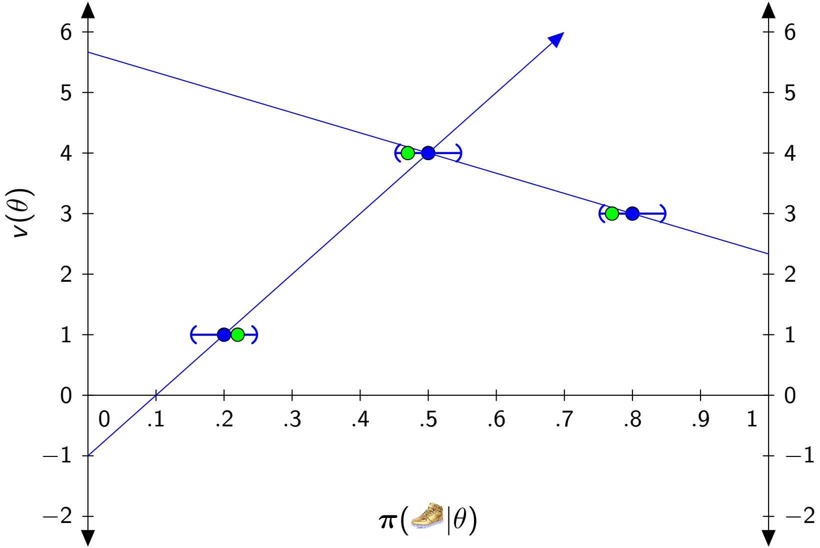
6
5
4
3
2
1
0
−1
−2
0
.1
.2
.3
.4
.5
.6
.7
.8
.9
1
$v(\theta)$
$\pi(\ |\theta)$

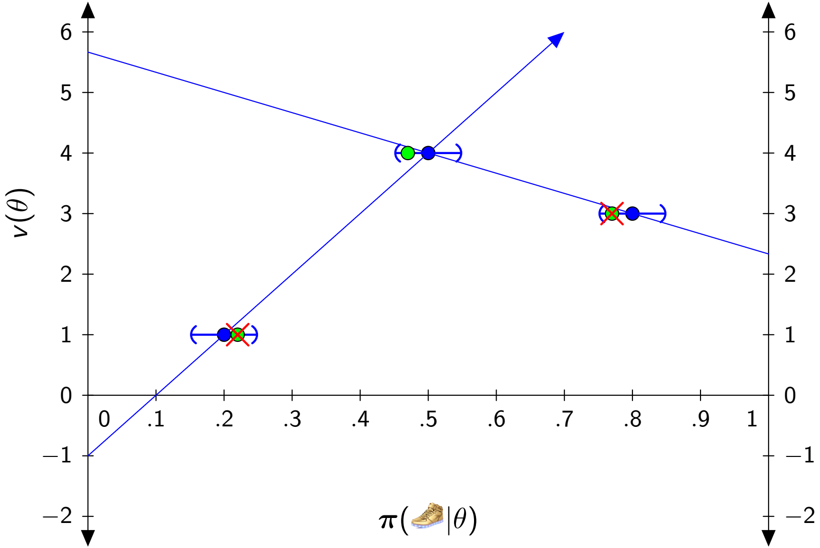
6
5
4
3
2
1
0
−1
−2
0 .1 .2 .3 .4 .5 .6 .7 .8 .9 1
$v(\theta)$
$\pi(\ |\theta)$

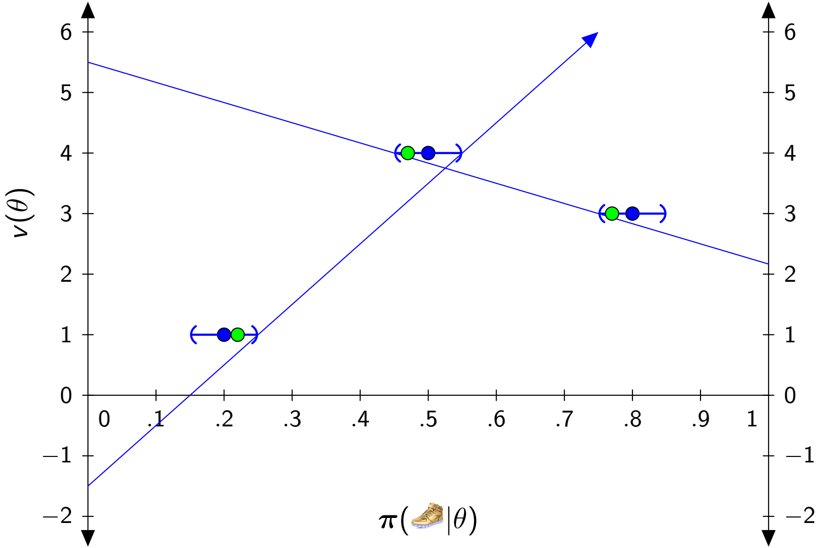
6
5
4
3
2
1
0
−1
−2
$v(\theta)$
0 .1 .2 .3 .4 .5 .6 .7 .8 .9 1
$\pi(\;|\theta)$

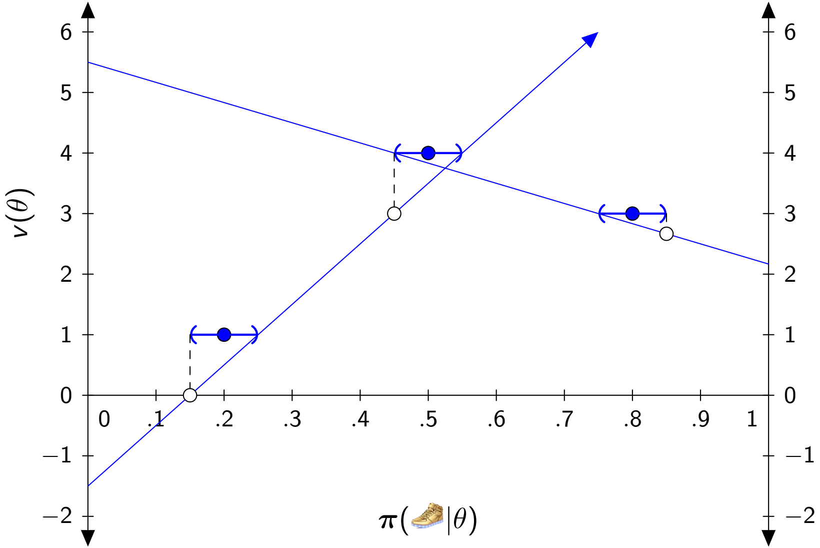

6
5
4
3
2
1
0
−1
−2
0 .1 .2 .3 .4 .5 .6 .7 .8 .9 1
$v(\theta)$
$\pi(\cdot|\theta)$

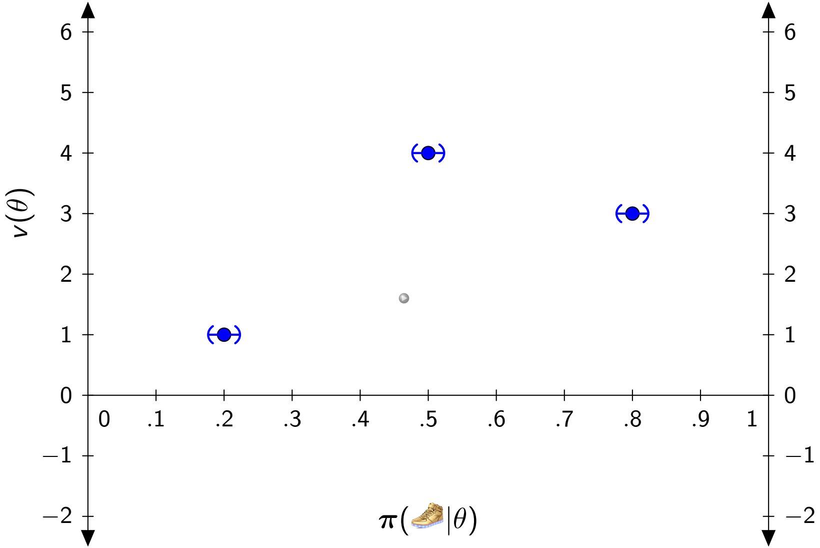
v(θ)
6
5
4
3
2
1
0
−1
−2
0
.1
.2
.3
.4
.5
.6
.7
.8
.9
1
π(|θ)

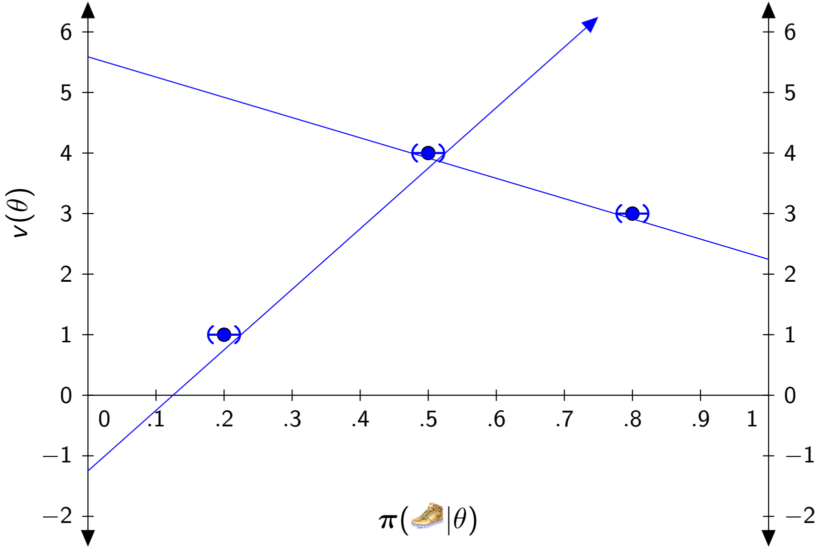

6
5
4
3
2
1
0
−1
−2
$v(\theta)$
0 .1 .2 .3 .4 .5 .6 .7 .8 .9 1
$\pi(\ |\theta)$

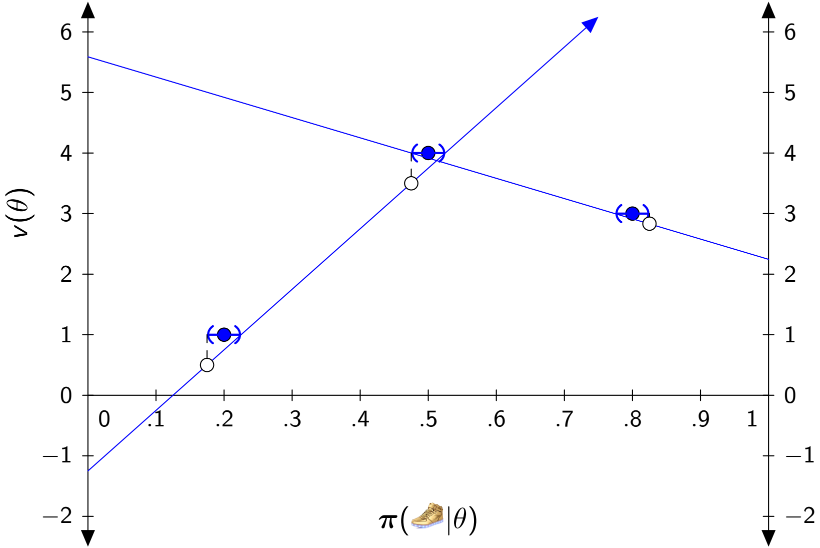
6
5
4
3
2
1
0
−1
−2
v(θ)
0
.1
.2
.3
.4
.5
.6
.7
.8
.9
1
π(|θ)

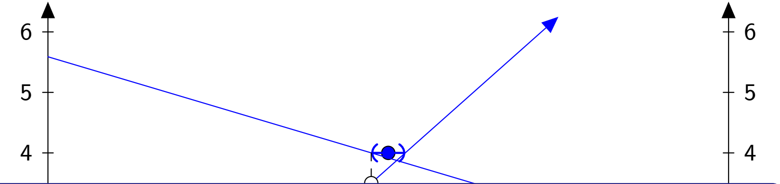

Theorem: Sufficient Correlation Implies Near Optimal Revenue

For any distribution $\boldsymbol{\pi}^*$ *that satisfies the FSE condition and given any positive constant* $k > 0$, **there exists** $\boldsymbol{\epsilon} > \mathbf{0}$ *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* $\mathbf{OPT}\,(\boldsymbol{\pi}^*) - \mathbf{k}$.

Return

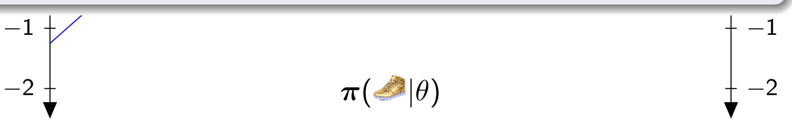

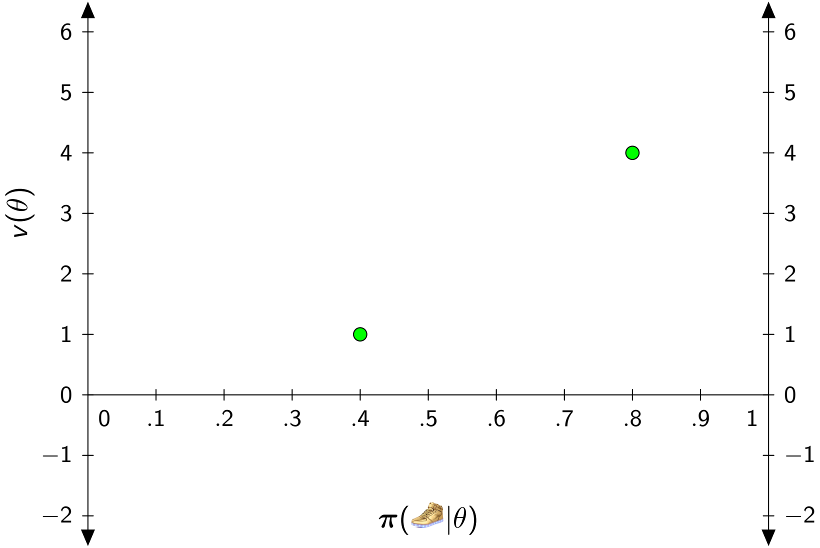
6
5
4
3
2
1
0
−1
−2
0 .1 .2 .3 .4 .5 .6 .7 .8 .9 1
$v(\theta)$
$\pi(\cdot|\theta)$

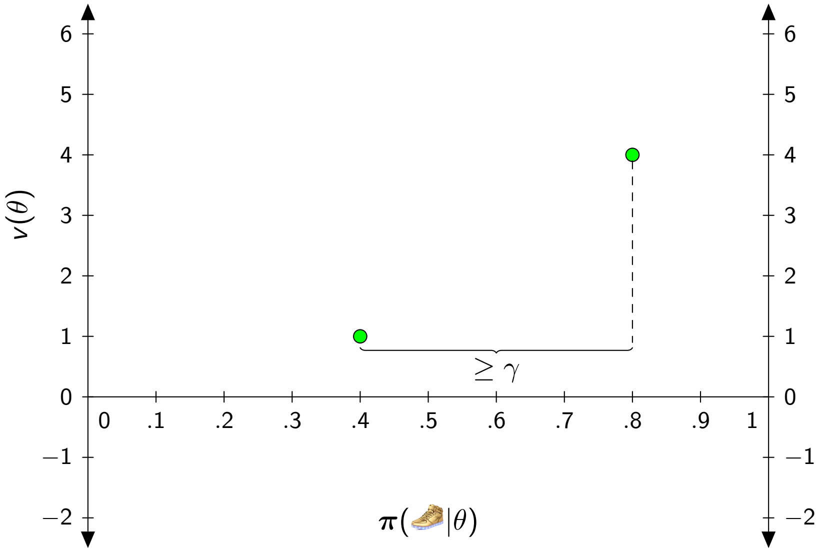
6
5
4
3
2
1
0
−1
−2
$v(\theta)$
0 .1 .2 .3 .4 .5 .6 .7 .8 .9 1
$\geq \gamma$
$\pi(\ |\theta)$

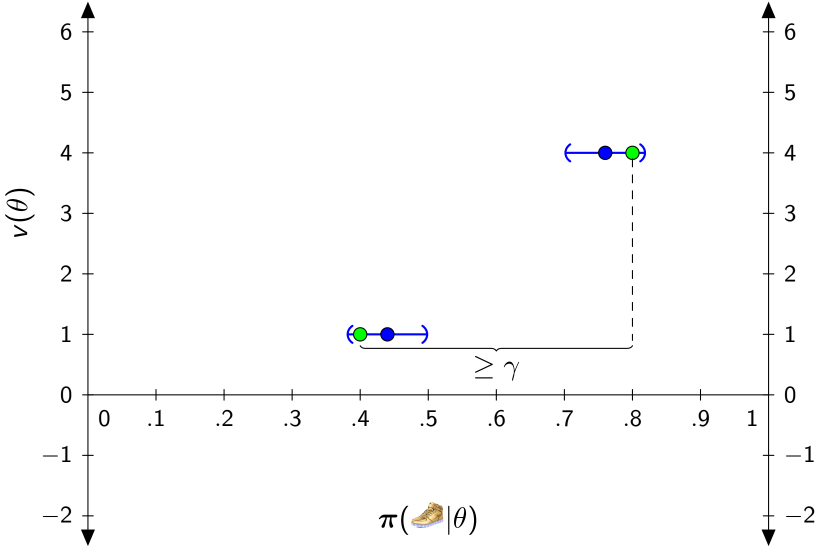

6
5
4
3
2
1
0
−1
−2
0
.1
.2
.3
.4
.5
.6
.7
.8
.9
1
$v(\theta)$
$\geq \gamma$
$\pi(\ |\theta)$

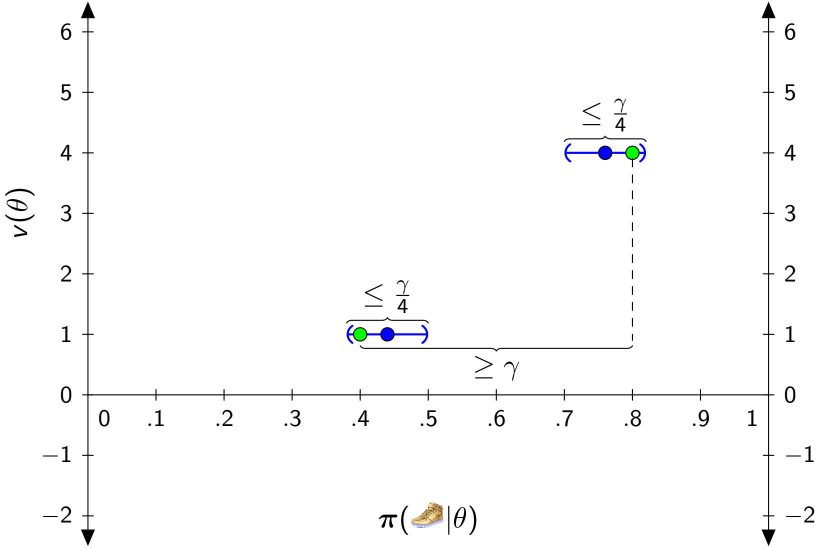

6
5
4
3
2
1
0
−1
−2
$v(\theta)$
$\leq \frac{\gamma}{4}$
$\leq \frac{\gamma}{4}$
$\geq \gamma$
0 .1 .2 .3 .4 .5 .6 .7 .8 .9 1
$\pi(\cdot|\theta)$

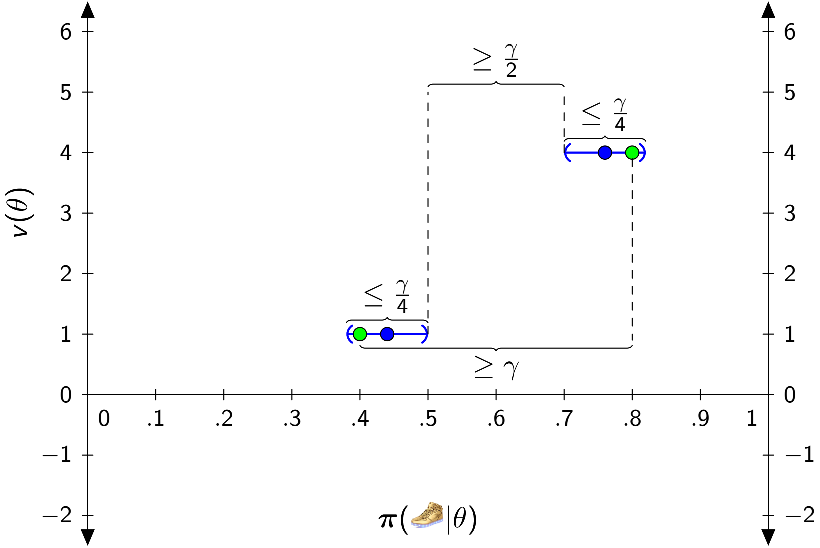

6
5
4
3
2
1
0
−1
−2
$v(\theta)$
$\geq \frac{\gamma}{2}$
$\leq \frac{\gamma}{4}$
$\leq \frac{\gamma}{4}$
$\geq \gamma$
0 .1 .2 .3 .4 .5 .6 .7 .8 .9 1
$\pi($
$|\theta)$

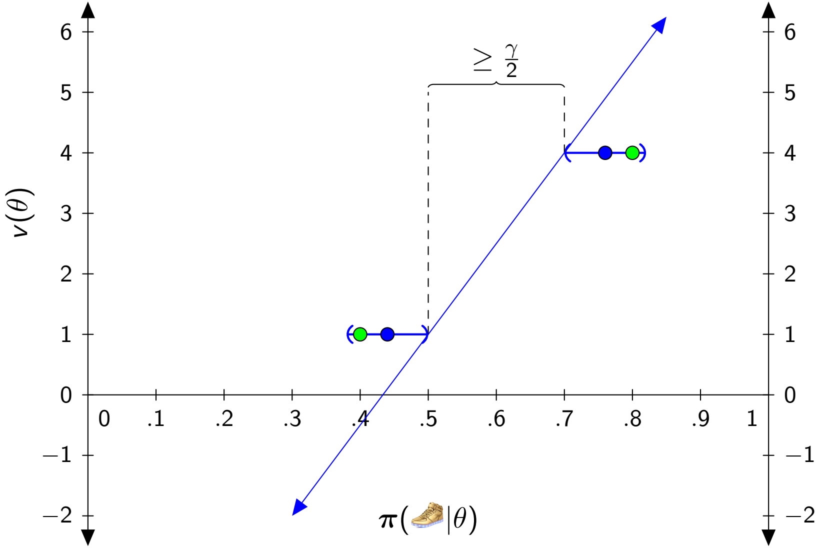
$\geq \frac{\gamma}{2}$
$v(\theta)$
$\pi(\text{👟}|\theta)$
6
5
4
3
2
1
0
−1
−2
0
.1
.2
.3
.4
.5
.6
.7
.8
.9
1

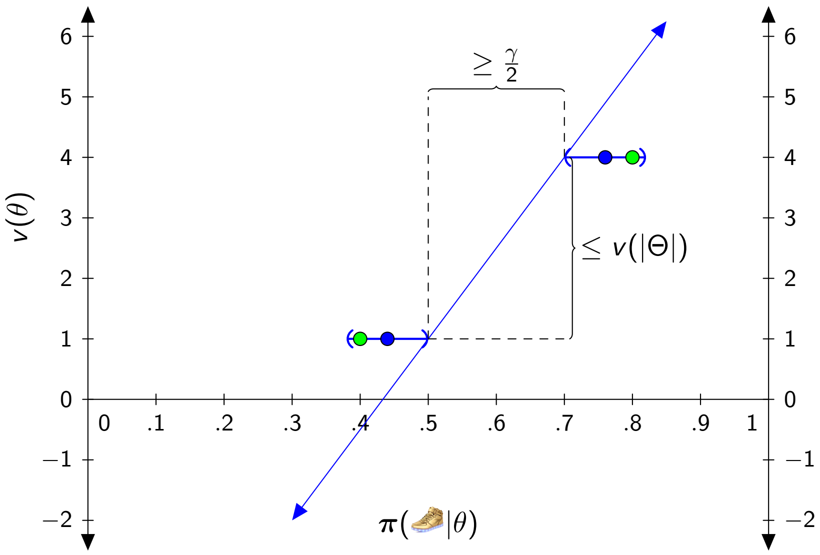
6
5
4
3
2
1
0
−1
−2
v(θ)
0 .1 .2 .3 .4 .5 .6 .7 .8 .9 1
≥ γ/2
≤ v(|Θ|)
π(|θ)

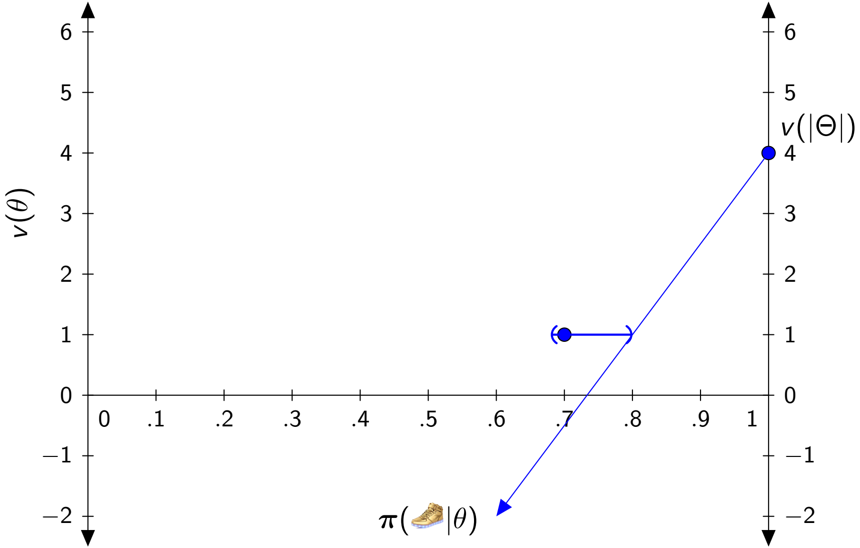
$v(\theta)$
6
5
4
3
2
1
0
−1
−2
0
.1
.2
.3
.4
.5
.6
.7
.8
.9
1
$v(|\Theta|)$
$\pi(\ |\theta)$

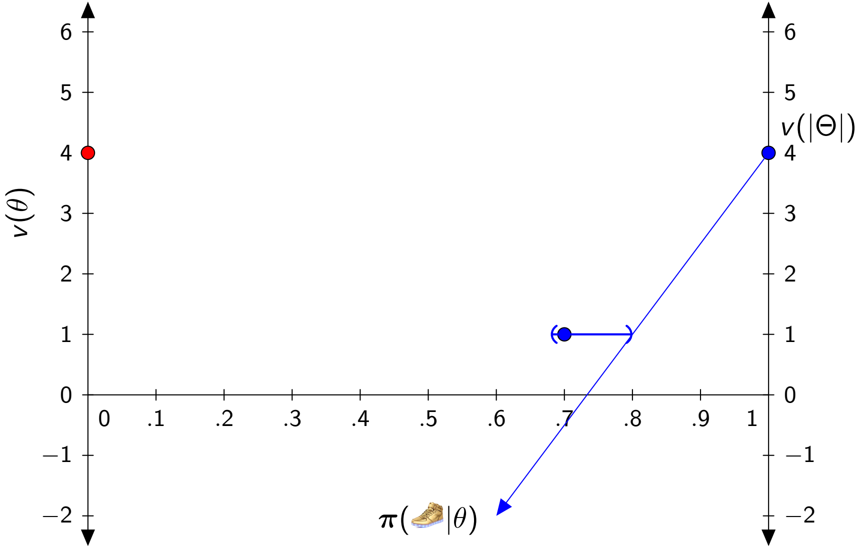
6
5
4
3
2
1
0
−1
−2
$v(\theta)$
0 .1 .2 .3 .4 .5 .6 .7 .8 .9 1
$v(|\Theta|)$
$\pi(\ |\theta)$

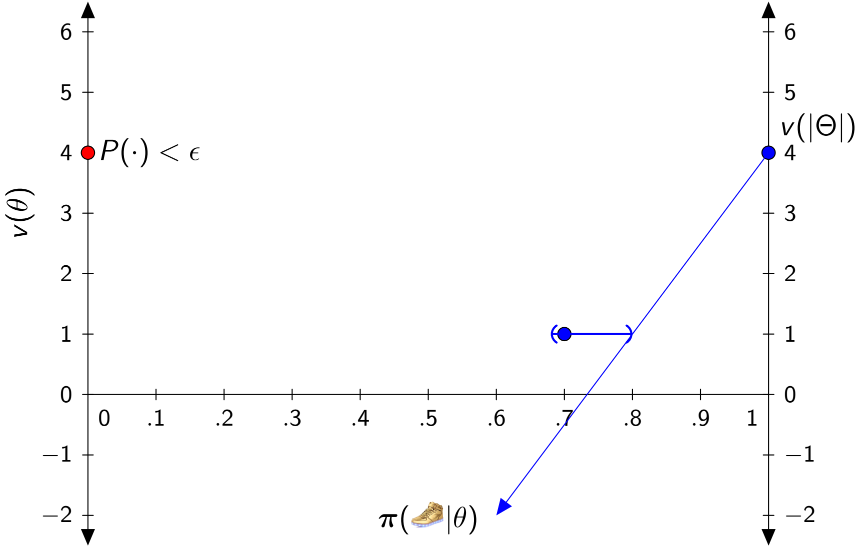

6
5
4
3
2
1
0
−1
−2
$v(\theta)$
$P(\cdot) < \epsilon$
$v(|\Theta|)$
0
.1
.2
.3
.4
.5
.6
.7
.8
.9
1
$\pi($ $|\theta)$

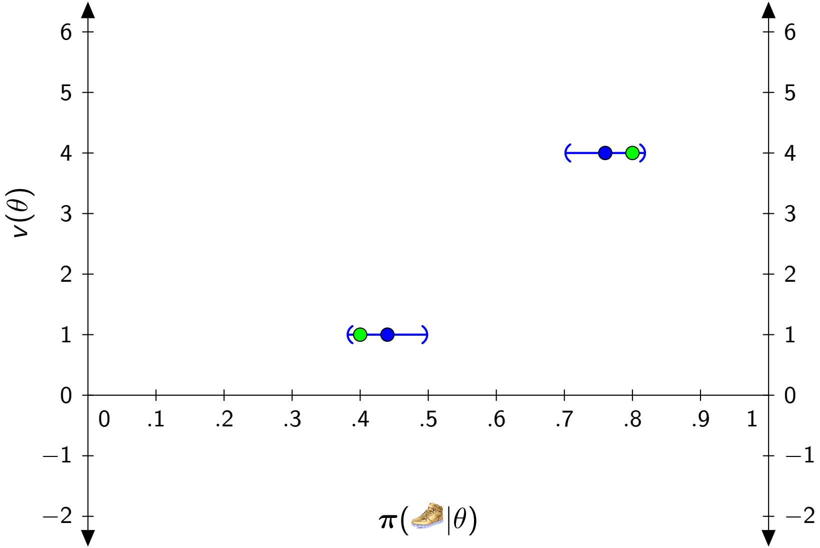
6
5
4
3
2
1
0
−1
−2
0 .1 .2 .3 .4 .5 .6 .7 .8 .9 1
v(θ)
π(|θ)

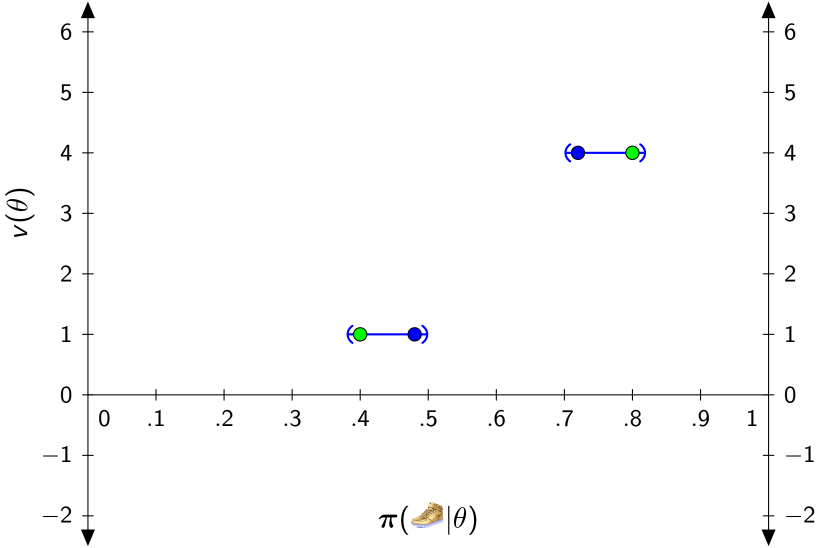
6
5
4
3
2
1
0
−1
−2
$v(\theta)$
0 .1 .2 .3 .4 .5 .6 .7 .8 .9 1
$\pi(\ |\theta)$

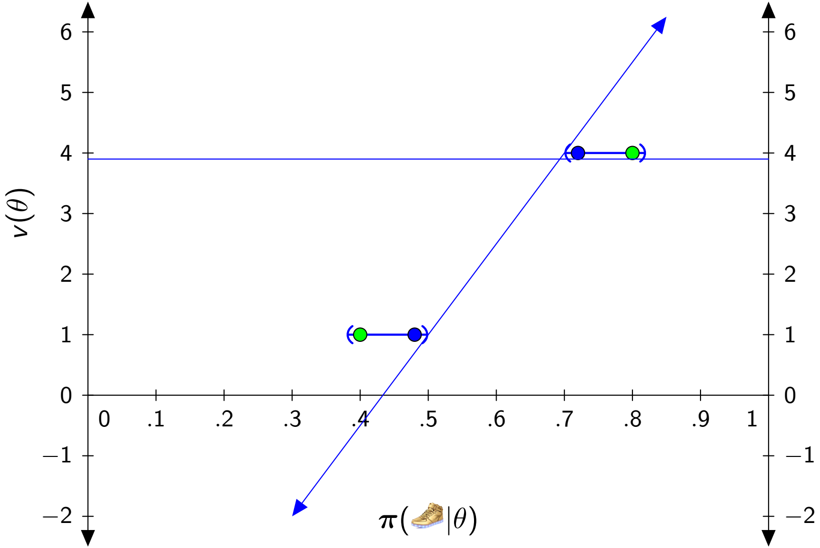
6
5
4
3
2
1
0
−1
−2
0
.1
.2
.3
.4
.5
.6
.7
.8
.9
1
$v(\theta)$
$\pi(\cdot|\theta)$

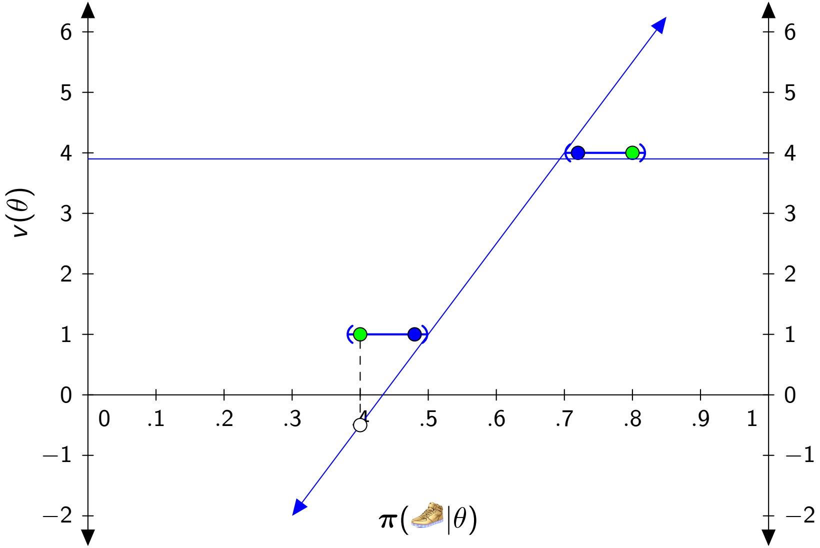
6
5
4
3
2
1
0
−1
−2
$v(\theta)$
0 .1 .2 .3 .4 .5 .6 .7 .8 .9 1
$\pi($ $|\theta)$

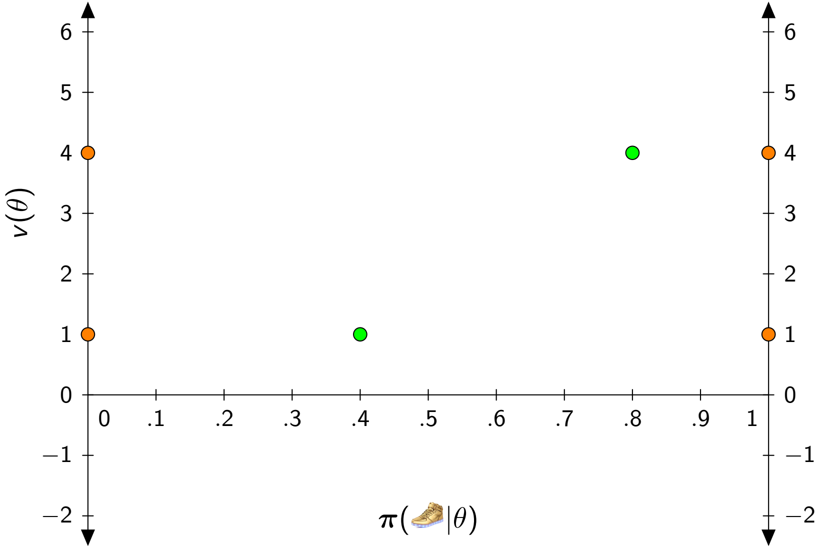
6
5
4
3
2
1
0
−1
−2
$v(\theta)$
0
.1
.2
.3
.4
.5
.6
.7
.8
.9
1
$\pi(\ |\theta)$

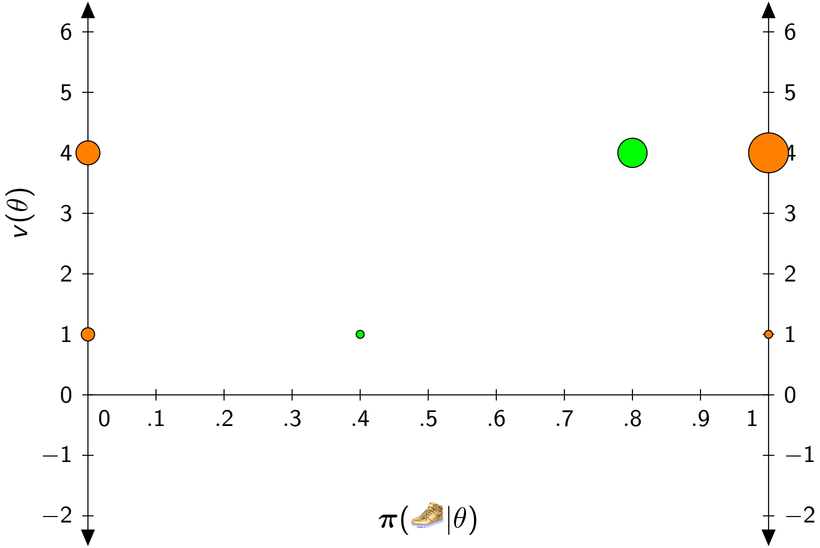
6
5
4
3
2
1
0
−1
−2
$v(\theta)$
0
.1
.2
.3
.4
.5
.6
.7
.8
.9
1
$\pi($ $|\theta)$

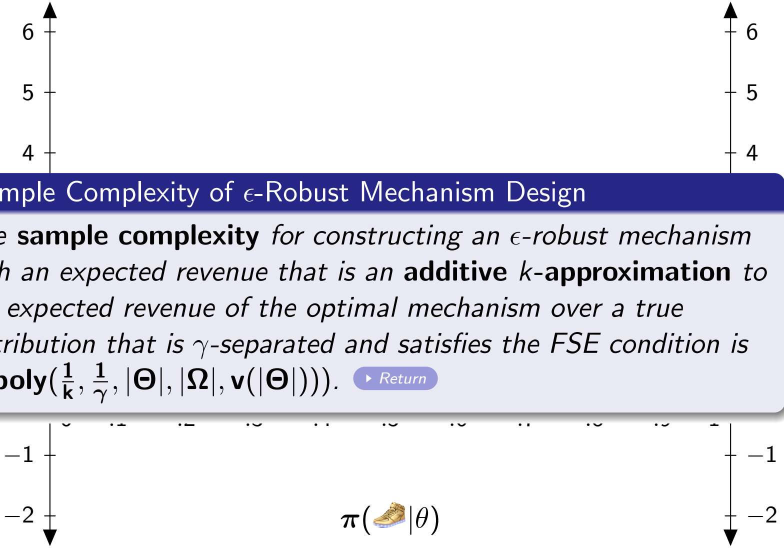

Sample Complexity of ϵ-Robust Mechanism Design

The **sample complexity** *for constructing an ϵ-robust mechanism with an expected revenue that is an* **additive k-approximation** *to the expected revenue of the optimal mechanism over a true distribution that is γ-separated and satisfies the FSE condition is* $\mathbf{O}(\mathbf{poly}(\frac{1}{\mathbf{k}}, \frac{1}{\gamma}, |\mathbf{\Theta}|, |\mathbf{\Omega}|, \mathbf{v}(|\mathbf{\Theta}|)))$.

Return

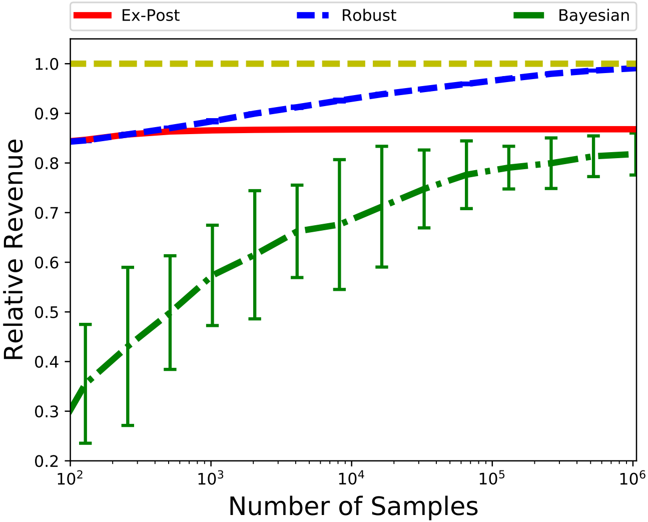
Ex-Post
Robust
Bayesian
Relative Revenue
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1.0
10^2
10^3
10^4
10^5
10^6
Number of Samples