Robust Bidding Policies: (Meaningfulness of Information for Bidding Decisions)

Saša Pekeč[†]

Decision Sciences, The Fuqua School of Business, Duke University

DIMACS/LAMSADE Meaningfulness Workshop
Paris, December 5, 2023.

[†]Joint with Chenxi Xu and V. Vinh Nguyen

Outline

Outline

Overview

- Focus on a single-bidder's optimization problem:
  bid decision that maximizes the bidder's payoff

Overview

- Focus on a single-bidder's optimization problem:
  bid decision that maximizes the bidder's payoff
- Bid decision depends on:
  - auction information: auction format (rules), items, information disclosure policies, etc.
  - bidder's own information, such as the objective, item value(s), etc.

Overview

- Focus on a single-bidder's optimization problem:
  bid decision that maximizes the bidder's payoff

- Bid decision depends on:
  - auction information: auction format (rules), items, information disclosure policies, etc.
  - bidder's own information, such as the objective, item value(s), etc.
  - information about other bidders' (rivals') objectives, valuations, behavior, etc.

Overview

- Focus on a single-bidder's optimization problem:
  bid decision that maximizes the bidder's payoff
- Bid decision depends on:
  - auction information: auction format (rules), items, information disclosure policies, etc.
  - bidder's own information, such as the objective, item value(s), etc.
  - information about other bidders' (rivals') objectives, valuations, behavior, etc.
- Research question:

  How does rivals' information impact the bidding policy?

  - is all rivals' information useful?
  - impact of (even ϵ) misspecification of rivals' information
  - impact of (distributional) assumptions about (uncertain) rivals' information

Bidding in auctions

Standard assumptions:

- Bidders are expected payoff maximizers.
- Same objective for each bidder.

Bidding in auctions

Standard assumptions:

- Bidders are expected payoff maximizers.
- Same objective for each bidder.

- What if objectives differ across bidders?
- What if uncertainty about rivals' objectives?

Bidding in auctions

Standard assumptions:

- Bidders are expected payoff maximizers.
- Same objective for each bidder.

- What if objectives differ across bidders?
- What if uncertainty about rivals' objectives?

Incentive compatibility: truthful reporting is a *dominant strategy*, i.e., maximizes bidder's payoff, regardless of rivals' actions.

- 2nd price (Vickrey) auction
- Uniform price multi-item auction (unit-demand, price set at the highest non-winning bid)
- Vickrey-Clarke-Groves (VCG) mechanisms

Bidding in auctions

Standard assumptions:

- Bidders are expected payoff maximizers.
- Same objective for each bidder.

- What if objectives differ across bidders?
- What if uncertainty about rivals' objectives?

Incentive compatibility: truthful reporting is a *dominant strategy*, i.e., maximizes bidder's payoff, regardless of rivals' actions.

- 2nd price (Vickrey) auction
- Uniform price multi-item auction (unit-demand, price set at the highest non-winning bid)
- Vickrey-Clarke-Groves (VCG) mechanisms

What about auctions that are not incentive compatible?

Example

First-price sealed-bid auction.

- $n = 2$ bidders and $m = 1$ item, IPV valuations $U[0,1]$
- Both bidders are **expected payoff** maximizers

Equilibrium bidding profile:

$$b_j^*(v) = v/2, \quad v \in [0,1], \quad j \in \{1,2\}$$

Example

First-price sealed-bid auction

- $n = 2$ bidders and $m = 1$ item, IPV valuations $U[0, 1]$
- Both bidders are **worst-case payoff** maximizers

Example

First-price sealed-bid auction

- $n = 2$ bidders and $m = 1$ item, IPV valuations $U[0, 1]$
- Both bidders are **worst-case payoff** maximizers

Truth-telling is an equilibrium

$$b_j^*(v) = v, \quad v \in [0, 1], \quad j \in \{1, 2\}$$

(multiple equilibria exist)

Example

First-price sealed bid auction

- $n = 2$ bidders and $m = 1$ item, IPV valuations $U[0,1]$
- Bidder 1 is a **worst-case payoff** maximizer
- Bidder 1 believes that Bidder 2 plays an expected payoff equilibrium strategy, i.e.,

$$b_2(v_2) = \frac{v_2}{2}, \quad v_2 \in [0,1]$$

Example

First-price sealed bid auction

- $n = 2$ bidders and $m = 1$ item, IPV valuations $U[0, 1]$
- Bidder 1 is a **worst-case payoff** maximizer
- Bidder 1 believes that Bidder 2 plays an expected payoff equilibrium strategy, i.e.,

$$b_2(v_2) = \frac{v_2}{2}, \quad v_2 \in [0, 1]$$

Bidder 1's best response is

$$b_1^*(v_1) = \min\left\{v_1, \frac{1}{2}\right\}, \quad v_1 \in [0, 1]$$

Example

First-price sealed bid auction

- $n = 2$ bidders and $m = 1$ item, IPV valuations $U[0, 1]$
- Bidder 1 is a **worst-case payoff** maximizer
- Bidder 1 believes that Bidder 2's **bid** b_2 is in $\mathcal{U}_{-1} = [0, d]$

Example

First-price sealed bid auction

- $n = 2$ bidders and $m = 1$ item, IPV valuations $U[0, 1]$
- Bidder 1 is a **worst-case payoff** maximizer
- Bidder 1 believes that Bidder 2's **bid** b_2 is in $\mathcal{U}_{-1} = [0, d]$

Bidder 1's best response is

$$b_1^*(v_1) = \min\{v_1, d\}, \quad v_1 \in [0, 1]$$

Related work

Incentive compatibility and market design

- Vickrey (1961), Clarke (1971), Groves (1973), Green & Laffont (1977), Rothkopf et al. (1990), Ausubel & Milgrom (2006), Roth (2015), Bichler (2017), Milgrom (2017,2019), etc.
- Edelman et al. (2007), Bichler & Goeree (2017), Cramton (2017), Karaenke et al. (2019), etc.

Use of robust optimization in analysis of auctions

- Ben-Tal & Nemirovski (2002), Bertsimas & Sim (2004), Bandi & Bertsimas (2014), etc.
- Bergemann & Shlag (2008), Kocyigit et al. (2018,2019), Carroll (2019), etc.

Core-selecting auctions

- Day & Raghavan (2007), Day & Milgrom (2008), Beck & Ott (2013), Ausubel & Baranov (2014,2017), Goeree & Lien (2016), etc.

Outline

Notation

- m indivisible items from the set $M = \{1, 2, \ldots, m\}$
- n bidders in $N = \{1, 2, \ldots, n\}$
- A *bundle* is a set S of items: $S \subseteq M$
- $v_j(S)$ is bidder j's valuation for bundle $S \subseteq M$
- $b_j(S)$ is bidder j's reported bid for bundle $S \subseteq M$
- Bidder j's payoff for obtaining bundle S_j at price p_j is $\pi_j = v_j(S_j) - p_j$ (quasi-linear preferences)
- Rivals' bid profile $b_{-j} = (b_1, b_2, \ldots, b_{j-1}, b_{j+1}, \ldots, b_n)$

Robust bidding problem

How should worst-case payoff maximizing bidder bid in auctions that are not incentive compatible?

Robust bidding problem

How should worst-case payoff maximizing bidder bid in auctions that are not incentive compatible?

- Bidder 1's feasible policy space $\mathcal{U}_1$

- Rivals' bids belong to an *uncertainty set* $\mathcal{U}_{-1}$

  Convex polytope: $\mathcal{U}_{-1} = \{b_{-1} \mid A\ b_{-1} \leq c\}$

  Box set: $\mathcal{U}_{-1} = \{b_{-1} \mid \underline{b}_j \leq b_j \leq \bar{b}_j, j = 2, 3, \ldots, n\}$

Robust bidding problem

How should worst-case payoff maximizing bidder bid in auctions that are not incentive compatible?

- Bidder 1's feasible policy space $\mathcal{U}_1$
- Rivals' bids belong to an *uncertainty set* $\mathcal{U}_{-1}$

Convex polytope: $\mathcal{U}_{-1} = \{b_{-1} \mid A\ b_{-1} \leq c\}$

Box set: $\mathcal{U}_{-1} = \{b_{-1} \mid \underline{b}_j \leq b_j \leq \bar{b}_j, j = 2, 3, \ldots, n\}$

Robust bidding problem

$$\pi_1^{MAXMIN} = \sup_{b_1 \in \mathcal{U}_1} \inf_{b_{-1} \in \mathcal{U}_{-1}} \pi_1(b_1, b_{-1})$$

Robust bidding problem

How should worst-case payoff maximizing bidder bid in auctions that are not incentive compatible?

- Bidder 1's feasible policy space $\mathcal{U}_1$
- Rivals' bids belong to an *uncertainty set* $\mathcal{U}_{-1}$

  Convex polytope: $\mathcal{U}_{-1} = \{b_{-1} \mid A\ b_{-1} \leq c\}$

  Box set: $\mathcal{U}_{-1} = \{b_{-1} \mid \underline{b}_j \leq b_j \leq \bar{b}_j, j = 2, 3, \ldots, n\}$

Robust bidding problem

$$\pi_1^{MAXMIN} = \sup_{b_1 \in \mathcal{U}_1} \inf_{b_{-1} \in \mathcal{U}_{-1}} \pi_1(b_1, b_{-1})$$

Let

$$\pi_1^{MINMAX} = \inf_{b_{-1} \in \mathcal{U}_{-1}} \sup_{b_1 \in \mathcal{U}_1} \pi_1(b_1, b_{-1})$$

Minimax inequality:

$$\pi_1^{MAXMIN} \leq \pi_1^{MINMAX}$$

Outline

Bidding in combinatorial auctions

- Combinatorial auctions: market-clearing of multiple heterogeneous items
- Multiple billions in value procured over last decade (spectrum, energy, pollution rights, real-estate, etc.)
- Combinatorial Clock Auction (Paul Milgrom, 2020 Nobel Prize)

Bidding in combinatorial auctions

- Combinatorial auctions: market-clearing of multiple heterogeneous items
- Multiple billions in value procured over last decade (spectrum, energy, pollution rights, real-estate, etc.)
- Combinatorial Clock Auction (Paul Milgrom, 2020 Nobel Prize)
- Combinatorial Clock Auction final phase: one-shot combinatorial auction
  - not incentive compatible, but
  - (the claim is that this is) mitigated by using core-selecting auctions,
  - (suggesting) truthful bidding the best bidder strategy from practical perspective, as it is not computationally/informationally burdensome and potential gains from strategizing might be elusive and could backfire.

Bidding in combinatorial auctions

- Combinatorial auctions: market-clearing of multiple heterogeneous items
- Multiple billions in value procured over last decade (spectrum, energy, pollution rights, real-estate, etc.)
- Combinatorial Clock Auction (Paul Milgrom, 2020 Nobel Prize)
- Combinatorial Clock Auction final phase: one-shot combinatorial auction
  - not incentive compatible, but
  - (the claim is that this is) mitigated by using core-selecting auctions,
  - (suggesting) truthful bidding the best bidder strategy from practical perspective, as it is not computationally/informationally burdensome and potential gains from strategizing might be elusive and could backfire.

Robust bidding is the best response to the rivals' truthful bidding in core-selecting auctions

Core-selecting auctions

We consider bidder-optimal core-selecting payment rule.
(Essentially, the closest point from the VCG outcome to the core.)

Core-selecting auctions

We consider bidder-optimal core-selecting payment rule.
(Essentially, the closest point from the VCG outcome to the core.)

Proposition

Let S_1 denote bidder 1's truthful allocation (i.e, the set of items S that bidder 1 receives by bidding truthfully).
Then

$$b_1^{PI}(S) = \begin{cases} 0 & \text{if } S \subsetneq S_1 \\ v_1(S_1) - \pi_1^{VCG} & \text{if } S_1 \subseteq S \subsetneq M \\ w_{b_{-1}}(N \setminus 1) & \text{if } S = M \end{cases}$$

is the optimal bidding policy for bidder 1.

- π_1^{VCG} is bidder 1's VCG payoff
- $w_{b_{-1}}(N \setminus 1)$ is the maximum surplus generated by allocating all items only to bidder 1's rivals

Suboptimality of truthful bidding

We consider bidder-optimal core-selecting payment rule.
(essentially, the closest point from the VCG outcome to the core)

Corollary

Assume bidders $2, \ldots, n$ bid truthfully. Then bidder 1 has a straightforward profitable deviation from bidding truthfully.

- If all rivals bid truthfully (as suggested "best" strategy in core selecting auctions), then bidding truthfully is not optimal.
  - For any $0 < \epsilon \leq \pi_1^{VCG}$, bidding $b_1(S_1) = v_1(S_1) - \epsilon$ is a profitable deviation.
  - Requires bidding $b_1(M) = w_{b_{-1}}(N \setminus 1)$.

Bidding under uncertainty: the single-minded bidder

- Bidder 1 is single-minded if it has positive valuation for a particular bundle S_1, i.e.,

$$v_1(S) = a > 0 \quad \text{if } S \supseteq S_1,$$
$$v_1(S) = 0 \quad \text{if } S \not\supseteq S_1.$$

Bidding under uncertainty: the single-minded bidder

- Bidder 1 is single-minded if it has positive valuation for a particular bundle S_1, i.e.,

$$v_1(S) = a > 0 \quad \text{if } S \supseteq S_1,$$
$$v_1(S) = 0 \quad \text{if } S \not\supseteq S_1.$$

Proposition

If bidder 1 is single-minded (and $v_1(S_1) > \bar{p}^{VCG}$) then a robust policy for bidder 1 is

$$b_1^{RO}(S) = v_1(S_1) - \min_{b_{-1} \in \mathcal{U}_{-1}} \pi_1^{VCG}, \quad \text{if } S_1 \subseteq S \subsetneq M,$$
$$b_1^{RO}(M) = v_1(S_1) - \min_{b_{-1} \in \mathcal{U}_{-1}} \pi_1^{VCG} + w_{b_{-1}}(N \setminus 1, M \setminus S_1),$$
$$b_1^{RO}(S) = 0, \quad \text{otherwise.}$$

Outline

L\L\G valuation structure

- $n = 3$ bidders and $m = 2$ homogeneous items
- Bidder 1: *local* bidder
- Rivals: Bidder 2 is *local* and bidder 3 is *global*

# items	v_1	b_1	b_2	b_3
1	a	x	b	0
2	a	y	b	c

L\L\G valuation structure

- $n = 3$ bidders and $m = 2$ homogeneous items
- Bidder 1: *local* bidder
- Rivals: Bidder 2 is *local* and bidder 3 is *global*

# items	v_1	b_1	b_2	b_3
1	a	x	b	0
2	a	y	b	c

- Box-type uncertainty set

$$\mathcal{U}_{-1} = \{(b_2, b_3) \mid \bar{b} - \epsilon_b \leq b \leq \bar{b} + \epsilon_b, \bar{c} - \epsilon_c \leq c \leq \bar{c} + \epsilon_c\}$$

L\L\G valuation structure

- $n = 3$ bidders and $m = 2$ homogeneous items
- Bidder 1: *local* bidder
- Rivals: Bidder 2 is *local* and bidder 3 is *global*

# items	v_1	b_1	b_2	b_3
1	a	x	b	0
2	a	y	b	c

- Box-type uncertainty set

$$\mathcal{U}_{-1} = \{(b_2, b_3) \mid \bar{b} - \epsilon_b \leq b \leq \bar{b} + \epsilon_b, \bar{c} - \epsilon_c \leq c \leq \bar{c} + \epsilon_c\}$$

- Let $\xi = \max_{b_{-1} \in \mathcal{U}_{-1}} (c - b)^+$
- Bidder 1's robust bidding policy:

$$b_1^{RO,1} = (\xi, \xi + \bar{b} - \epsilon_b)$$

L\L\G valuation structure: numerical example

- $n = 3$ bidders and $m = 2$ homogeneous items
- Bidder 1: *local* bidder
- Rivals: Bidder 2 is *local* and bidder 3 is *global*

# items	v_1	b_1	b_2	b_3
1	10	x	b	0
2	10	y	b	c

- Box-type uncertainty set

$$\mathcal{U}_{-1} = \{(b_2, b_3) \mid \quad 7 \leq b \leq 13, \quad 7 \leq c \leq 13\}$$

- Bidder 1 robust bidding policy:

$$b_1^{RO} = (6, 13)$$

Robust bidding vs. truthful bidding

- Under robust policy $b_1^{RO} = (6, 13)$, bidder 1's payoff is point-wise greater than her truthful payoff

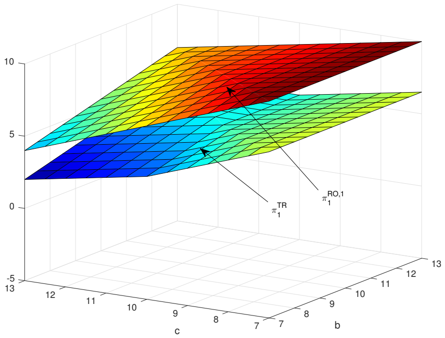

Comparison of b_1^{RO} and b_1^{TR}

Robust bidding vs. (misspecified) perfect info bidding

Suppose bidder 1 believes its rivals will bid b and c

- Bidder 1 best response is denoted by b_1^{PI}

Robust bidding vs. (misspecified) perfect info bidding

Suppose bidder 1 believes its rivals will bid b and c

- Bidder 1 best response is denoted by b_1^{PI}

If there is any uncertainty about the rivals' actual bidding profile:

$$\mathcal{U}_{-1} = [b - \epsilon, b + \epsilon] \times [c - \epsilon, c + \epsilon],$$

then one can compare performance of the (misspecified) perfect information bidding b_1^{PI} and robust bidding b_1^{RO}.

Robust bidding vs. (misspecified) perfect info bidding

Suppose bidder 1 believes its rivals will bid b and c

- Bidder 1 best response is denoted by b_1^{PI}

If there is any uncertainty about the rivals' actual bidding profile:

$$\mathcal{U}_{-1} = [b - \epsilon, b + \epsilon] \times [c - \epsilon, c + \epsilon],$$

then one can compare performance of the (misspecified) perfect information bidding b_1^{PI} and robust bidding b_1^{RO}.

Numerical example:

- Set $b = 10$, $c = 10$
- Then $b_1^{PI} = (0, 10)$
- Let $\epsilon = 3$
- Then $b_1^{RO} = (6, 13)$

Robust bidding vs. (misspecified) perfect info bidding

Let $b = c = 10$, and let $\epsilon = 3$.

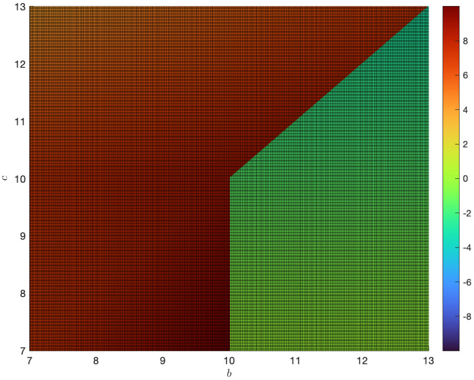

Payoff difference for $b_1^{RO} = (6, 13)$ and $b_1^{PI} = (0, 10)$

Robust bidding vs. expected-payoff maximization

- When b and c are uniformly distributed on $[7, 13]$, the robust bidding and expected-payoff maximization policies are the same
- We also consider non-uniform distributions.
  e.g., $f_b \uparrow, f_c \uparrow$ has linearly increasing marginal densities:

$$f_b(y) = f_c(y) = (y-7)/18, \quad 7 \leq y \leq 13.$$

Robust bidding vs. expected-payoff maximization

- When b and c are uniformly distributed on $[7, 13]$, the robust bidding and expected-payoff maximization policies are the same
- We also consider non-uniform distributions.
  e.g., $f_b \uparrow, f_c \uparrow$ has linearly increasing marginal densities:

$$f_b(y) = f_c(y) = (y - 7)/18, \quad 7 \leq y \leq 13.$$

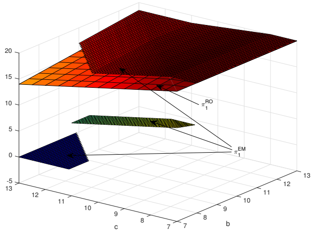

Comparison of b_1^{RO} and b_1^{EM}

Robust bidding policy performance

$\bar{b} = \bar{c} = 10$ and $\epsilon_b = \epsilon_c = 3$

Distr.	EM	RO	TR	PI	PI ($\epsilon'_b = \epsilon'_c = 1$)
$f_b, f_c \uparrow$	7.9092	7.6010	4.7253	4.7284	7.1695
		(96.10%)	(59.74%)	(59.78%)	(90.65%)
$f_b, f_c \downarrow$	8.5904	8.5904	5.2267	2.8548	4.7150
		(100%)	(60.84%)	(33.23%)	(54.89%)
$f_b \uparrow, f_c \downarrow$	9.1212	8.9027	5.5274	7.1266	8.5371
		(97.60%)	(60.60%)	(78.13%)	(93.60%)
$f_b \downarrow, f_c \uparrow$	6.8877	6.8877	4.0236	1.2946	3.1104
		(100%)	(58.42%)	(18.80%)	(45.16%)
$\mathcal{N}$	8.5508	8.2207	4.9186	3.8648	7.6708
		(96.14%)	(57.52%)	(45.20%)	(89.71%)
Triangular	8.3014	8.1500	4.9000	3.8986	6.8839
		(98.18%)	(59.03%)	(46.96%)	(82.92%)

Robust bidding policy performance

$\bar{b} = 10$, $\bar{c} = 13.5$, and $\epsilon_b = \epsilon_c = 3$

Distr.	EM	RO	TR	PI	PI ($\epsilon'_b = \epsilon'_c = 1$)
$f_b, f_c \uparrow$	4.9728	4.4871	3.2362	3.3582	4.6680
		(90.23%)	(65.08%)	(67.53%)	(93.87%)
$f_b, f_c \downarrow$	5.4876	5.4870	3.2361	0.7910	2.0499
		(99.99%)	(58.97%)	(14.41%)	(37.36%)
$f_b \uparrow, f_c \downarrow$	6.7200	6.3850	4.1342	5.5605	6.4263
		(95.01%)	(61.52%)	(82.75%)	(95.63%)
$f_b \downarrow, f_c \uparrow$	3.5130	3.4996	2.2487	0.3855	1.1891
		(99.62%)	(64.01%)	(10.97%)	(33.85%)
$\mathcal{N}$	5.3573	4.9989	3.2489	2.0920	4.6624
		(93.31%)	(60.64%)	(39.05%)	(87.03%)
Triangular	5.1933	4.9950	3.2450	2.2286	4.1309
		(96.18%)	(62.48%)	(42.91%)	(79.54%)

Summary

- Focus on single-bidder optimization problem: bidding policies that maximize worst-case payoff.
- Bidder uncertainty is modeled via robust optimization framework (uncertainty set)
  - belief-free re rivals' valuations
  - belief-free re rivals' objectives
- Minimax (in)equality is the key argument in the proofs.

Summary

- Focus on single-bidder optimization problem: bidding policies that maximize worst-case payoff.
- Bidder uncertainty is modeled via robust optimization framework (uncertainty set)
  - belief-free re rivals' valuations
  - belief-free re rivals' objectives
- Minimax (in)equality is the key argument in the proofs.
- Robust bidding in challenging settings, such as core-selecting auctions:
  - easy to determine bidding policy
  - bypasses challenges with processing rivals' information, such as misspecification, distributional assumptions, objective, behavior
  - outperforms truthful bidding
  - outperforms misspecified perfect information setting

Thank you!

pekec@duke.edu