Estimating Approximate Incentive Compatibility

Maria-Florina Balcan, Tuomas Sandholm,
and **Ellen Vitercik**
Berkeley (EECS) → Stanford (MS&E and CS)

Working paper, preliminary version in
Conference on Economics and Computation (EC)

Incentive compatibility (IC)

Agents maximize utility by reporting type truthfully

Fundamental concept in mechanism design

Many real-world mechanisms are **not** IC

Discriminatory auctions

Multi-unit variant of first-price auction

Not incentive compatible

Used to sell treasury bills since 1929

Discriminatory auctions

Multi-unit variant of first-price auction

Not incentive compatible

Used to sell treasury bills since 1929
and electricity in the UK

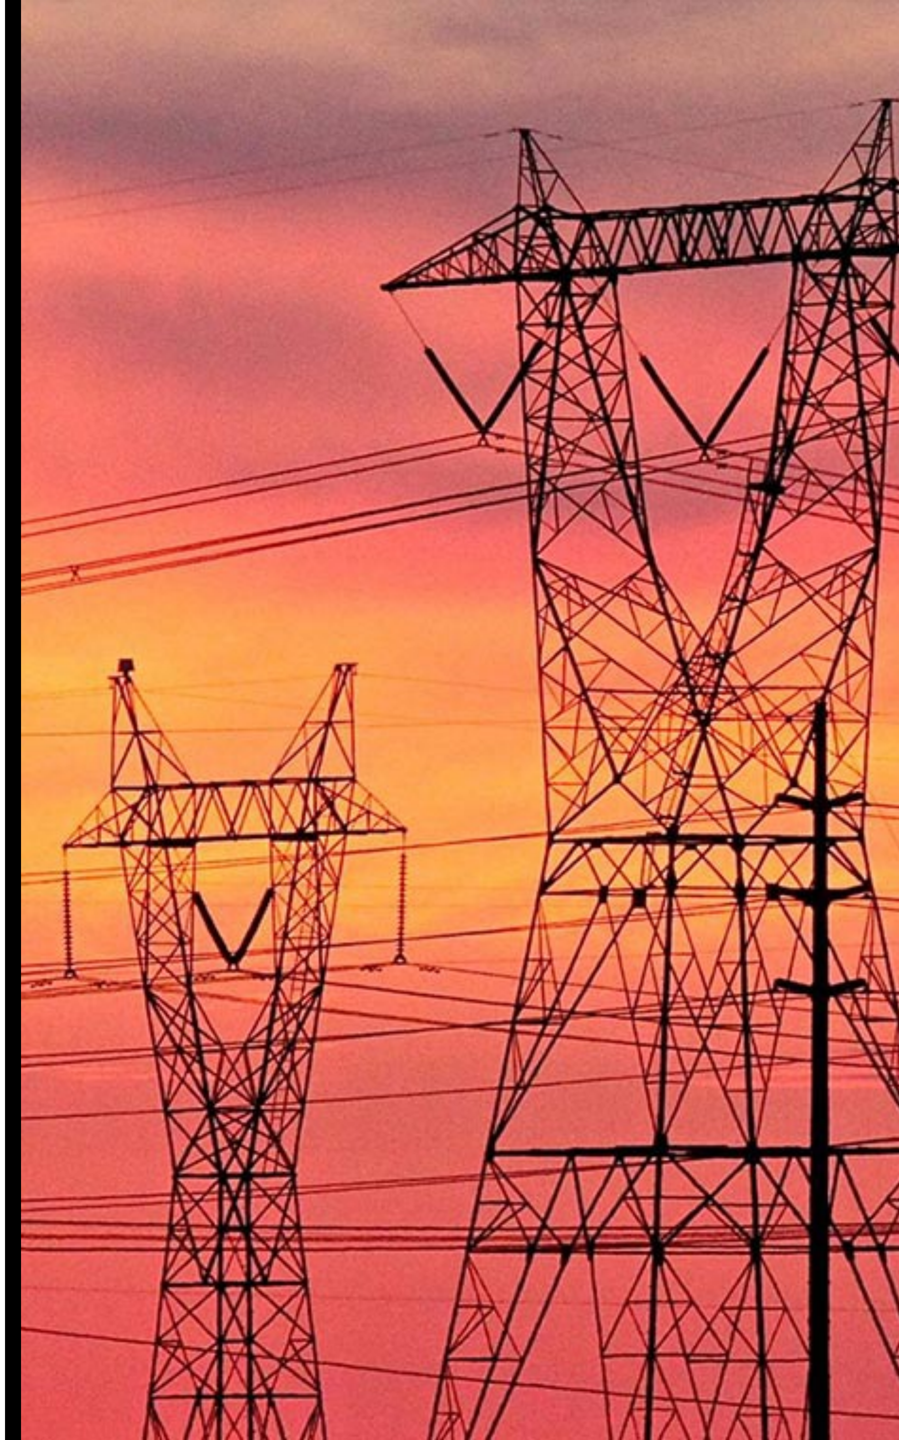

Generalized 2^{nd}-price

Used for sponsored search

Not incentive compatible

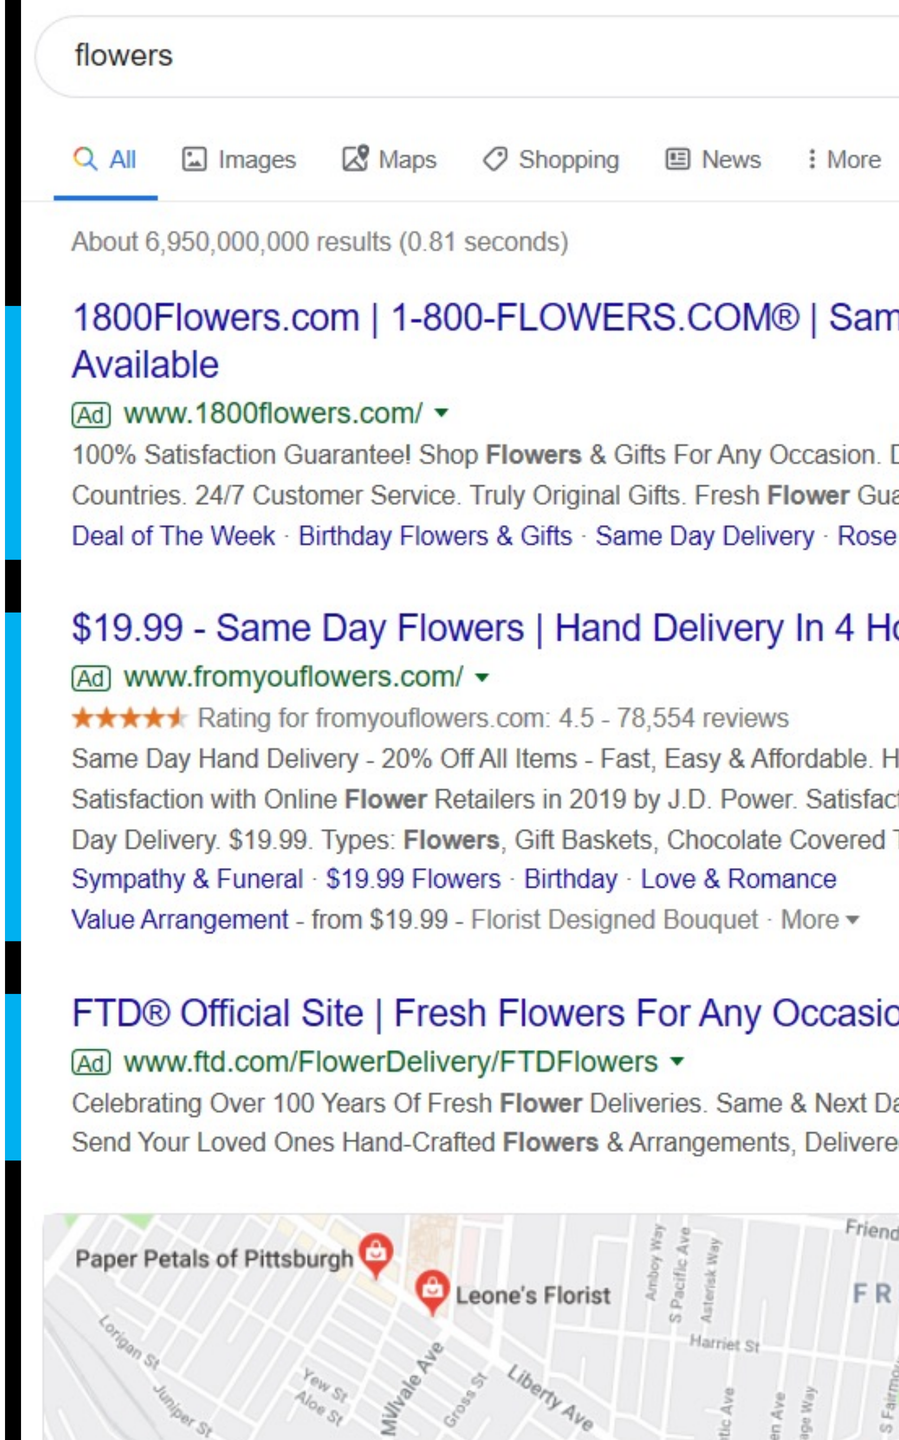

Multi-item, multi-agent auctions

Nearly all fielded combinatorial auctions

(such as sourcing auctions)

aren't incentive compatible

Why aren't real-world auctions IC?

Rules are **easier** to explain

Bids used to tune **future** parameters

Might leak **private** values

Agents not **risk** neutral

Approximate incentive compatibility

Auction is **γ-IC** when for each agent i:

If everyone except agent i is truthful,

she can only increase exp. utility by γ when lies about type

Kothari, Parkes, Suri, EC'03; Archer, Papadimitriou, Talwar, Tardos, Internet Mathematics '04; Conitzer and Sandholm, IJCAI'07; Dekel, Fischer, Procaccia, JCSS'10; Lubin, Parkes, Current Science '12; Mennle and Seuken, EC'14; Dütting, Fischer, Jirapinyo, Lai, Lubin, Parkes TEAC'15; Azevedo, Budish, Review of Economic Studies '18; Feng, Narasimhan, Parkes, AAMAS'18; Golowich, Narasimhan, Parkes, IJCAI'18; Dütting, Feng, Narasimhan, Parkes, Ravindranath, ICML'19

Approximate incentive compatibility

Auction is **γ-IC** when for each agent i:

If everyone except agent i is truthful,

she can only increase exp. utility by γ when lies about type

Overriding goal: Given **samples** from dist. over agents' types,

estimate IC approximation factor (γ) using samples

Complements literature on sample-based **revenue maximization**

Likhodedov, Sandholm, AAAI'04, '05; Balcan, Blum, Hartline, Mansour, FOCS'05; Elkind, SODA'07; Cole, Roughgarden, STOC'14; Mohri, Medina, ICML'14; Huang, Mansour, Roughgarden, EC'15; Morgenstern, Roughgarden, NeurIPS'15, COLT'16; Roughgarden, Schrijvers, EC'16; Devanur, Huang, Psomas, STOC'16; Balcan, Sandholm, **V.**, NeurIPS'16; Gonczarowski, **Nisan**, STOC'17; Cai, Daskalakis, FOCS'17; Balcan, Sandholm, **V.**, EC'18; …

Why estimate IC approximation factor?

Some mechanisms might have terrible **worst case** IC apx factor,
but are (nearly) IC for distribution over agents' types at hand

Use mechanism might have discarded as non-IC,
remaining optimistic that agents will be truthful

Why estimate IC approximation factor?

In mechanism design via machine learning:

Add constraint requiring this estimate be small

[Feng, Narasimhan, Parkes, AAMAS'18; Golowich, Narasimhan, Parkes, IJCAI'18; Dütting, Feng, Narasimhan, Parkes, Ravindranath, ICML'19]

Is resulting mechanism (nearly) IC?

Background

n agents with types in $[0,1]^M$

Standard assumption:

Agents' types drawn from probability distribution $(\boldsymbol{t}_1, \ldots, \boldsymbol{t}_n) \sim \mathcal{D}$

In this talk, $\mathcal{D}$ is product distribution: $\boldsymbol{t}_i \sim \mathcal{D}_i$ and $\mathcal{D} = \mathcal{D}_1 \times \cdots \times \mathcal{D}_n$

$\mathcal{D}_{-i} = \times_{j \neq i} \mathcal{D}_j$

Incentive compatibility

Auction is **incentive compatible** (IC) if for any agent:

In expectation over other's types,

utility maximized by reporting type truthfully,

so long as others also truthful

Utility of agent i: $u(t_i, \tilde{t}_i, \boldsymbol{t}_{-i})$

Incentive compatibility

Auction is **incentive compatible** (IC) if for any agent:

In expectation over other's types,

utility maximized by reporting type truthfully,

so long as others also truthful

Utility of agent i: $u(t_i, \tilde{t}_i, \boldsymbol{t}_{-i})$

True type

Incentive compatibility

Auction is **incentive compatible** (IC) if for any agent:
In expectation over other's types,
utility maximized by reporting type truthfully,
so long as others also truthful

Utility of agent i: $u(t_i, \tilde{t}_i, \boldsymbol{t}_{-i})$

Reported type

Incentive compatibility

Auction is **incentive compatible** (IC) if for any agent:

In expectation over other's types,

utility maximized by reporting type truthfully,

so long as others also truthful

Utility of agent i: $u(t_i, \tilde{t}_i, \boldsymbol{t}_{-i})$

Others' types

Incentive compatibility

Auction is **incentive compatible** (IC) if for any agent:

In expectation over other's types,

utility maximized by reporting type truthfully,

so long as others also truthful

Mechanism is IC if for any agent i and any $t_i, \tilde{t}_i$,

$$\mathbb{E}_{\boldsymbol{t}_{-i}}[u(t_i, t_i, \boldsymbol{t}_{-i})] \geq \mathbb{E}_{\boldsymbol{t}_{-i}}[u(t_i, \tilde{t}_i, \boldsymbol{t}_{-i})]$$

Utility from truthful report

Utility from strategic report

"Ex-interim" IC

Approximate incentive compatibility

Auction is **γ-IC** if for any agent i and any $t_i, \tilde{t}_i$,

$$\mathbb{E}_{\boldsymbol{t}_{-i}}[u(t_i, t_i, \boldsymbol{t}_{-i})] \geq \mathbb{E}_{\boldsymbol{t}_{-i}}[u(t_i, \tilde{t}_i, \boldsymbol{t}_{-i})] - \gamma$$

Utility from truthful report

Utility from strategic report

Kothari, Parkes, Suri, EC'03; Archer, Papadimitriou, Talwar, Tardos, Internet Mathematics '04; Conitzer and Sandholm, IJCAI'07; Dekel, Fischer, Procaccia, JCSS'10; Lubin, Parkes, Current Science '12; Mennle and Seuken, EC'14; Dütting, Fischer, Jirapinyo, Lai, Lubin, Parkes TEAC'15; Azevedo, Budish, Review of Economic Studies '18; Feng, Narasimhan, Parkes, AAMAS'18; Golowich, Narasimhan, Parkes, IJCAI'18; Dütting, Feng, Narasimhan, Parkes, Ravindranath, ICML'19

Our goal: Estimate IC approximation factor (γ) using samples

Our estimate

Our estimate (first try):

Maximum utility agent i can gain by misreporting her type,

on average over samples $\boldsymbol{t}_{-i}^{(1)}, \ldots, \boldsymbol{t}_{-i}^{(N)} \sim \mathcal{D}_{-i}$:

$$\max_{t_i, \tilde{t}_i \in \mathbb{R}^M} \left\{ \frac{1}{N} \sum_{j=1}^{N} u\left(t_i, \tilde{t}_i, \boldsymbol{t}_{-i}^{(j)}\right) - u\left(t_i, t_i, \boldsymbol{t}_{-i}^{(j)}\right) \right\}$$

Utility from strategic report

Utility from truthful report

Our estimate

Our estimate (first try):

Maximum utility agent i can gain by misreporting her type, on average over samples $\boldsymbol{t}_{-i}^{(1)}, \ldots, \boldsymbol{t}_{-i}^{(N)} \sim \mathcal{D}_{-i}$:

$$\max_{t_i, \tilde{t}_i \in \mathbb{R}^M} \left\{ \frac{1}{N} \sum_{j=1}^{N} u\left(t_i, \tilde{t}_i, \boldsymbol{t}_{-i}^{(j)}\right) - u\left(t_i, t_i, \boldsymbol{t}_{-i}^{(j)}\right) \right\}$$

Not convex and many discontinuities

Our estimate

Our estimate $\widehat{\gamma}$:

Maximum utility agent i can gain by misreporting her type,
on average over samples $\boldsymbol{t}_{-i}^{(1)}, \dots, \boldsymbol{t}_{-i}^{(N)} \sim \mathcal{D}_{-i}$,
if true & reported types from uniform grid $\mathcal{G}$

$$\widehat{\gamma} = \max_{t_i, \tilde{t}_i \in \mathcal{G}} \left\{ \frac{1}{N} \sum_{j=1}^{N} u\left(t_i, \tilde{t}_i, \boldsymbol{t}_{-i}^{(j)}\right) - u\left(t_i, t_i, \boldsymbol{t}_{-i}^{(j)}\right) \right\}$$

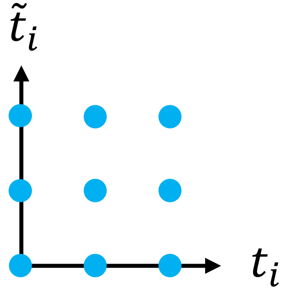

Our estimate

Our estimate $\hat{\gamma}$:

Maximum utility agent i can gain by misreporting her type,

on average over samples $\boldsymbol{t}_{-i}^{(1)}, \ldots, \boldsymbol{t}_{-i}^{(N)} \sim \mathcal{D}_{-i}$,

if true & reported types from uniform grid $\mathcal{G}$

$$\hat{\gamma} = \max_{t_i, \tilde{t}_i \in \mathcal{G}} \left\{ \frac{1}{N} \sum_{j=1}^{N} u\left(t_i, \tilde{t}_i, \boldsymbol{t}_{-i}^{(j)}\right) - u\left(t_i, t_i, \boldsymbol{t}_{-i}^{(j)}\right) \right\}$$

Challenge:

Might miss pairs of true & reported types with large utility gains

Our estimate

Our estimate $\hat{\gamma}$:

Maximum utility agent i can gain by misreporting her type,

on average over samples $\boldsymbol{t}_{-i}^{(1)}, \dots, \boldsymbol{t}_{-i}^{(N)} \sim \mathcal{D}_{-i}$,

if true & reported types from uniform grid $\mathcal{G}$

$$\hat{\gamma} = \max_{t_i, \tilde{t}_i \in \mathcal{G}} \left\{ \frac{1}{N} \sum_{j=1}^{N} u\left(t_i, \tilde{t}_i, \boldsymbol{t}_{-i}^{(j)}\right) - u\left(t_i, t_i, \boldsymbol{t}_{-i}^{(j)}\right) \right\}$$

Key question:

$|\gamma - \hat{\gamma}| \leq$ **?**

Uniform grid

Challenge:

Utility functions are volatile

First-price auction:

- Highest bidder wins
- Pays highest bid

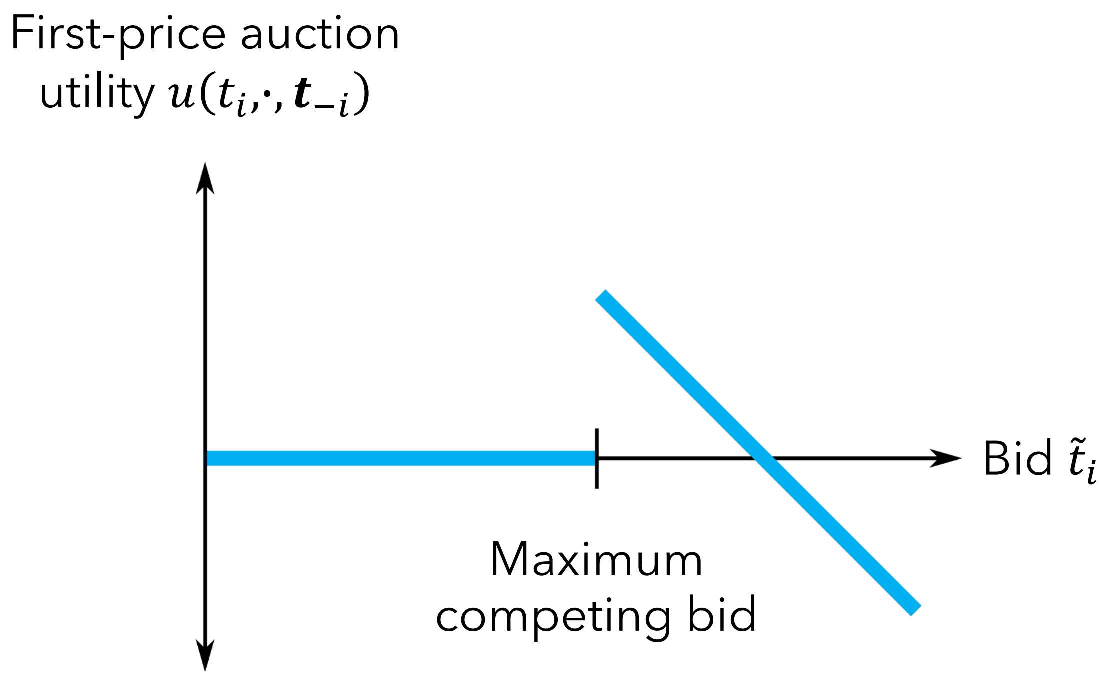

Uniform grid

Challenge:

Utility functions are volatile, especially on average over samples

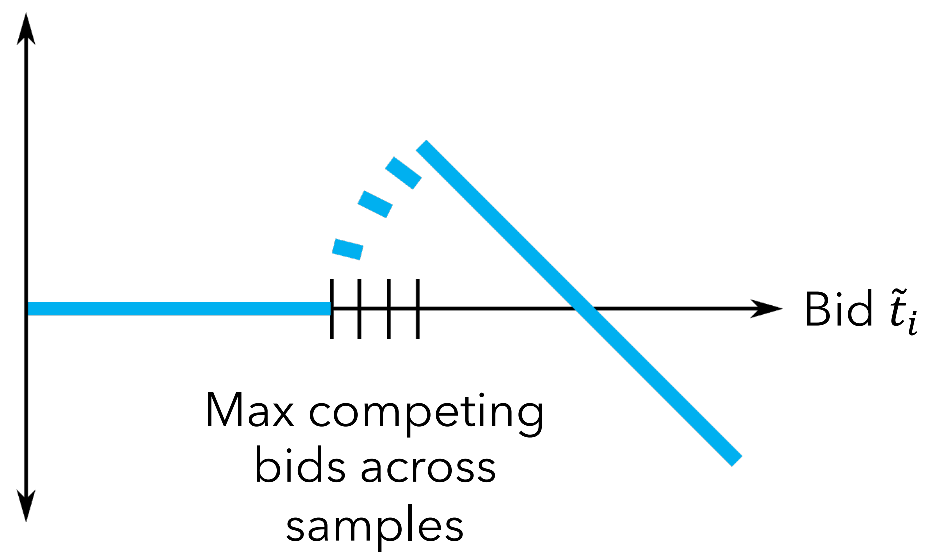

Uniform grid

Coarse discretization can lead to poor utility estimation

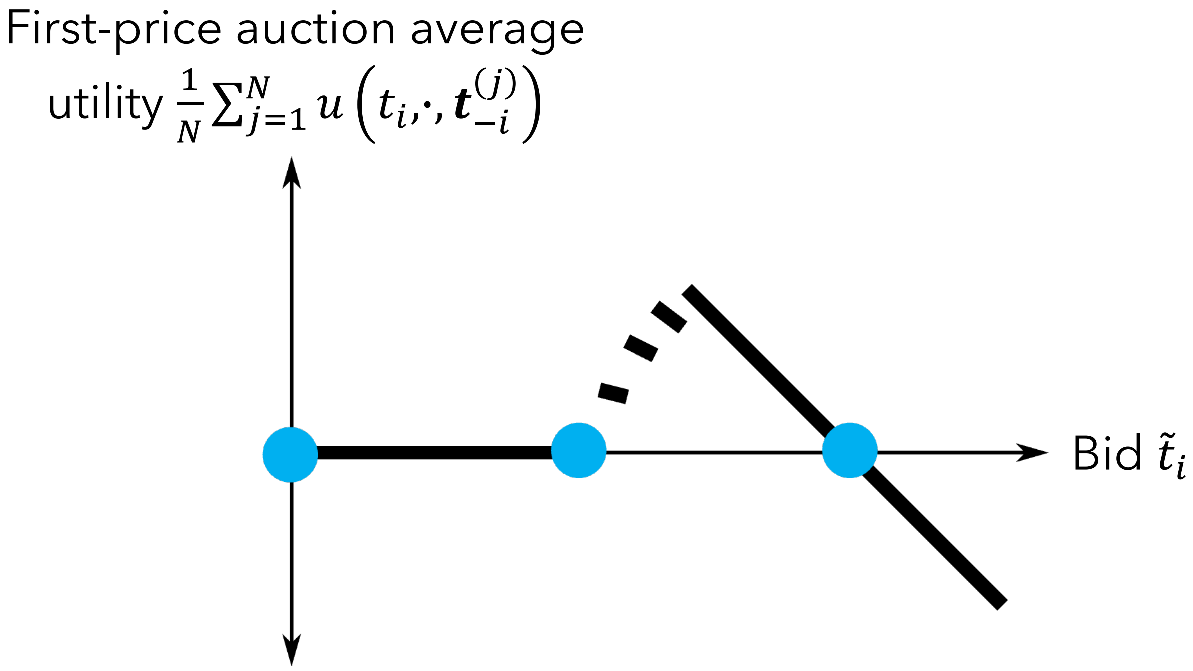

When is the distribution “nice” enough to use a grid?

Dispersion

Functions $f_1, \ldots, f_N : \mathbb{R}^M \to \mathbb{R}$ are **(w, k)-dispersed** if:

Every w-ball contains discontinuities of $\leq k$ functions

[Balcan, Dick, **V.**, FOCS'18]

Plot $\sum f_i$:

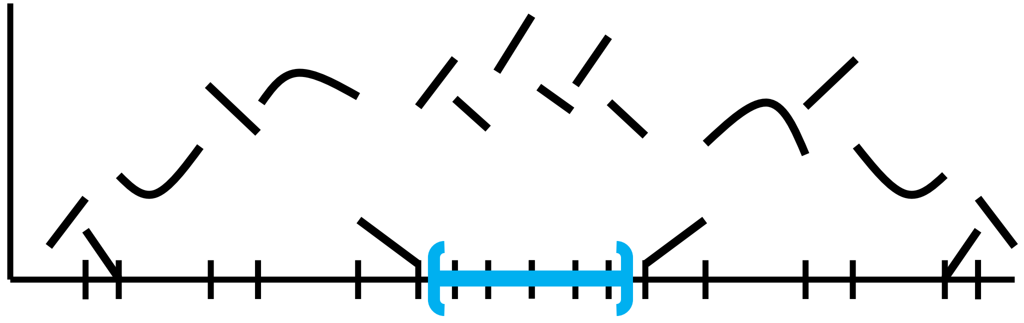

Not dispersed

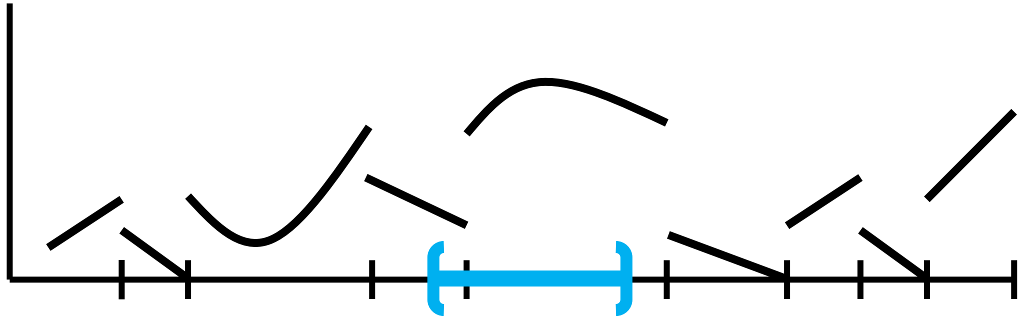

Dispersed

Dispersed utility functions

Definition. $\boldsymbol{t}_{-i}^{(1)}, \dots, \boldsymbol{t}_{-i}^{(N)}$ **induce L-Lipschitz (w, k)-dispersion** if:

1. For any t_i, $u\left(t_i, \cdot, \boldsymbol{t}_{-i}^{(1)}\right), \dots, u\left(t_i, \cdot, \boldsymbol{t}_{-i}^{(N)}\right)$ are:

Utility as a function of the bid $\tilde{t}_i$

Dispersed utility functions

Definition. $\boldsymbol{t}_{-i}^{(1)}, \dots, \boldsymbol{t}_{-i}^{(N)}$ **induce L-Lipschitz (w, k)-dispersion** if:

1. For any t_i, $u\left(t_i, \cdot, \boldsymbol{t}_{-i}^{(1)}\right), \dots, u\left(t_i, \cdot, \boldsymbol{t}_{-i}^{(N)}\right)$ are:
   Piecewise L-Lipschitz and (w, k)-dispersed

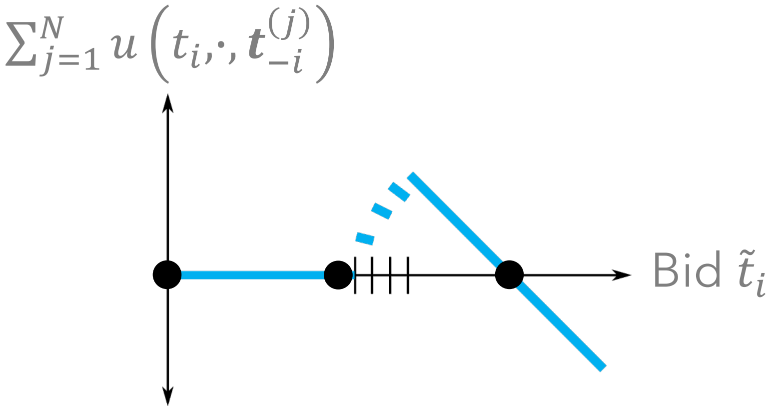

Not dispersed

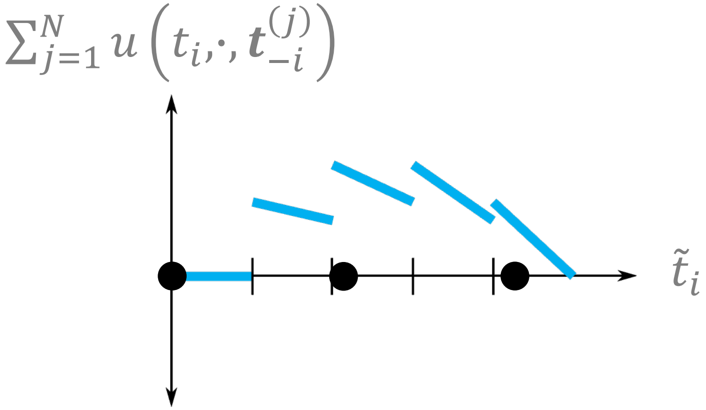

Dispersed

Dispersed utility functions

Definition. $\boldsymbol{t}_{-i}^{(1)}, \ldots, \boldsymbol{t}_{-i}^{(N)}$ **induce** **L-Lipschitz (w, k)-dispersion** if:

1. For any t_i, $u\left(t_i, \cdot, \boldsymbol{t}_{-i}^{(1)}\right), \ldots, u\left(t_i, \cdot, \boldsymbol{t}_{-i}^{(N)}\right)$ are:

   Piecewise L-Lipschitz and (w, k)-dispersed

2. For any $\tilde{t}_i$, $u\left(\cdot, \tilde{t}_i, \boldsymbol{t}_{-i}^{(1)}\right), \ldots, u\left(\cdot, \tilde{t}_i, \boldsymbol{t}_{-i}^{(N)}\right)$ are:

Utility as a function of the value t_i

Dispersed utility functions

Definition. $\boldsymbol{t}_{-i}^{(1)}, \dots, \boldsymbol{t}_{-i}^{(N)}$ **induce L-Lipschitz (w, k)-dispersion** if:

1. For any t_i, $u\left(t_i, \cdot, \boldsymbol{t}_{-i}^{(1)}\right), \dots, u\left(t_i, \cdot, \boldsymbol{t}_{-i}^{(N)}\right)$ are:
   Piecewise L-Lipschitz and (w, k)-dispersed
2. For any $\tilde{t}_i$, $u\left(\cdot, \tilde{t}_i, \boldsymbol{t}_{-i}^{(1)}\right), \dots, u\left(\cdot, \tilde{t}_i, \boldsymbol{t}_{-i}^{(N)}\right)$ are:
   Piecewise L-Lipschitz and (w, k)-dispersed

- Prove that WHP, for all **infinitely-many** function sequences: dispersion holds for "good" values of w and k
  - Show discontinuities are shared across function sequences

Guarantees

Our estimate $\hat{\gamma}$:

Maximum utility agent i can gain by misreporting her type,

on average over samples $\boldsymbol{t}_{-i}^{(1)}, \ldots, \boldsymbol{t}_{-i}^{(N)} \sim \mathcal{D}_{-i}$,

if true & reported types from uniform grid $\mathcal{G}$

Theorem:

If WHP, for all i, $\boldsymbol{t}_{-i}^{(1)}, \ldots, \boldsymbol{t}_{-i}^{(N)}$ induce L-Lipschitz (w, k)-dispersion

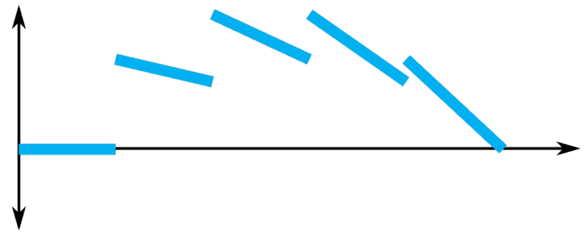

Guarantees

Our estimate $\hat{\gamma}$:

Maximum utility agent i can gain by misreporting her type,

on average over samples $\boldsymbol{t}_{-i}^{(1)}, \ldots, \boldsymbol{t}_{-i}^{(N)} \sim \mathcal{D}_{-i}$,

if true & reported types from uniform grid $\mathcal{G}$

Theorem:

If WHP, for all i, $\boldsymbol{t}_{-i}^{(1)}, \ldots, \boldsymbol{t}_{-i}^{(N)}$ induce L-Lipschitz (w, k)-dispersion

$\Rightarrow$ Can estimate using w-grid

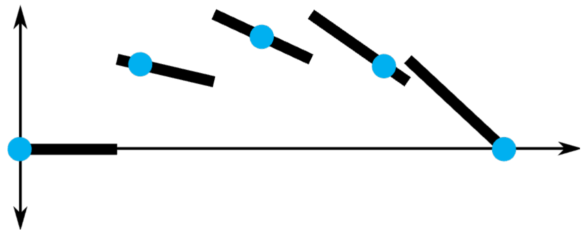

Guarantees

Theorem:

If WHP, for all i, $\boldsymbol{t}_{-i}^{(1)}, \ldots, \boldsymbol{t}_{-i}^{(N)}$ induce L-Lipschitz (w, k)-dispersion

⇒ Can estimate using w-grid

Estimation error: WHP, $|\hat{\gamma} - \gamma| = \tilde{O}\left(Lw + \frac{k}{N} + \sqrt{\frac{d}{N}}\right)$

d = standard ML measure of utility functions' intrinsic complexity

Guarantees

Theorem:

If WHP, for all i, $\boldsymbol{t}_{-i}^{(1)}, \ldots, \boldsymbol{t}_{-i}^{(N)}$ induce L-Lipschitz (w, k)-dispersion

$\Rightarrow$ Can estimate using w-grid

Estimation error: WHP, $|\hat{\gamma} - \gamma| = \tilde{O}\left(\boldsymbol{Lw} + \frac{\boldsymbol{k}}{\boldsymbol{N}} + \sqrt{\frac{d}{N}}\right)$

Proof idea:

- If snap types to grid, average utility only changes by $\leq \boldsymbol{Lw} + \frac{\boldsymbol{k}}{\boldsymbol{N}}$

Guarantees

Theorem:

If WHP, for all i, $\boldsymbol{t}_{-i}^{(1)}, \ldots, \boldsymbol{t}_{-i}^{(N)}$ induce L-Lipschitz (w, k)-dispersion

⇒ Can estimate using w-grid

Estimation error: WHP, $|\hat{\gamma} - \gamma| = \tilde{O}\left(Lw + \frac{k}{N} + \sqrt{\frac{d}{N}}\right)$

Proof idea:

- If snap types to grid, average utility only changes by $\leq Lw + \frac{k}{N}$
- $\sqrt{\frac{d}{N}}$ additional error incurred from sampling

Guarantees

Theorem:

If WHP, for all i, $\boldsymbol{t}_{-i}^{(1)}, \ldots, \boldsymbol{t}_{-i}^{(N)}$ induce L-Lipschitz (w, k)-dispersion

⇒ Can estimate using w-grid

Estimation error: WHP, $|\hat{\gamma} - \gamma| = \tilde{O}\left(Lw + \frac{k}{N} + \sqrt{\frac{d}{N}}\right)$

When $w = O\left(\frac{1}{\sqrt{N}}\right), k = O(\sqrt{N})$:

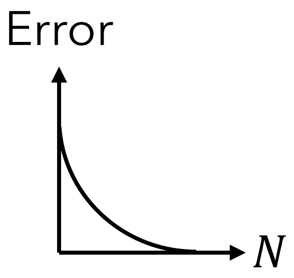

We prove these (w, k) values hold when distribution is **nice**

Applications

When does dispersion hold?

$[0, \kappa]$ = range of density functions defining agents' type distributions

First-price auction

Error: $|\hat{\gamma} - \gamma| = \tilde{O}\left(\frac{(\#\text{agents}) + \kappa^{-1}}{\sqrt{(\#\text{samples})}}\right)$

Also analyze **combinatorial** first-price auctions

Applications

When does dispersion hold?

$[0, \kappa]$ = range of density functions defining agents' type distributions

Generalized second-price auction

Error: $|\hat{\gamma} - \gamma| = \tilde{O}\left(\frac{(\#\text{agents})^{3/2} + \kappa^{-1}}{\sqrt{(\#\text{samples})}}\right)$

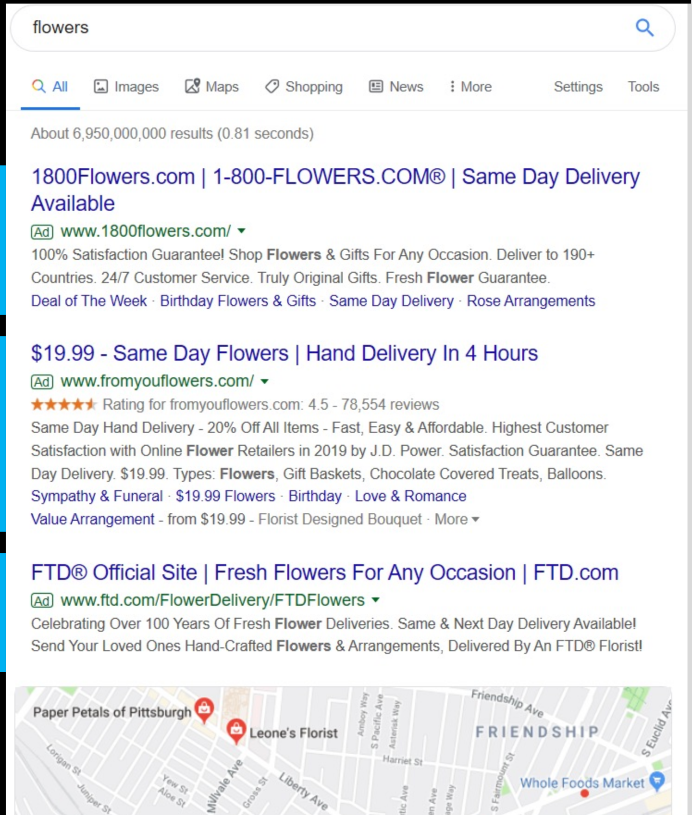

Applications

When does dispersion hold?

$[0, \kappa]$ = range of density functions defining agents' type distributions

Discriminatory and uniform price auctions

Generalization of first-price auction to multi-unit settings

Error: $|\hat{\gamma} - \gamma| = \tilde{O}\left(\frac{(\#\text{agents})(\#\text{units})^2 + \kappa^{-1}}{\sqrt{(\#\text{samples})}}\right)$

Conclusion

- Provide techniques for estimating how far mechanism is from IC
- Introduce empirical variant of approximate IC
- Bound estimate's error using *dispersion*
- Guarantees for:
  - First-price (combinatorial) auction
  - Generalized second-price auction
  - Discriminatory auction
  - Uniform price auction
  - Second-price auction under spiteful agents

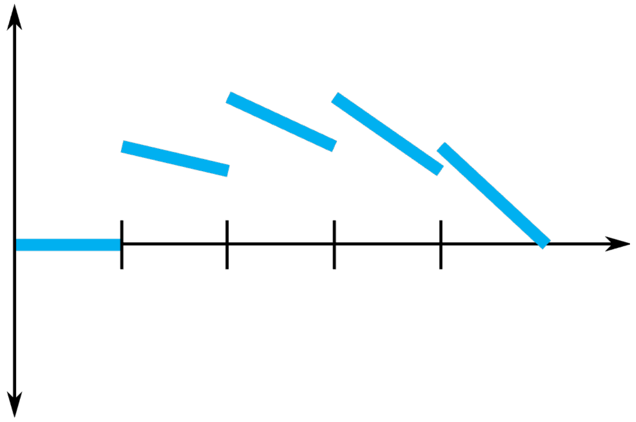

Future directions

What if samples strategically manipulated?

Also applies to literature on revenue maximization via machine learning

Likhodedov, Sandholm, AAAI'04, '05; Elkind, SODA'07; Cole, Roughgarden, STOC'14; Mohri, Medina, ICML'14; Huang, Mansour, Roughgarden, EC'15; Morgenstern, Roughgarden, NeurIPS'15, COLT'16; Devanur, Huang, Psomas, STOC'16; Balcan, Sandholm, **V.**, NeurIPS'16; Gonczarowski, Nisan, STOC'17; Cai, Daskalakis, FOCS'17; Balcan, Sandholm, **V.**, EC'18; …

What about when report space not real-valued?

E.g., school choice mechanisms: report preference order over schools

Estimating Approximate Incentive Compatibility

Maria-Florina Balcan, Tuomas Sandholm,
and **Ellen Vitercik**
Berkeley (EECS) → Stanford (MS&E and CS)

Working paper, preliminary version in
Conference on Economics and Computation (EC)