Estimating Approximate Incentive Compatibility

Maria-Florina Balcan, Tuomas Sandholm, and Ellen Vitercik

Appeared in EC 2019

Incentive compatibility (IC)

Fundamental concept in mechanism design
Buyers maximize their utilities by bidding truthfully

Many real-world mechanisms are not IC:
Discriminatory auctions
- Multi-unit variant of first-price auction
- Used to sell US treasury bills and UK electricity

Generalized second price auction
- Used for sponsored search

First-price auction
- Display ad exchanges may be transitioning to FP

Most fielded combinatorial auctions
- For example, sourcing auctions

Why not?
- Might be expensive for buyers to compute true values
- Rules are often easier to explain
- Bids used to tune future auction
- Auction might leak the bidders' private values
- Bidders might not be risk neutral

Approximate IC

Mechanism is **γ-IC** when for each bidder i:
If everyone except bidder i is truthful,
she can only increase utility by γ if she bids strategically

Defined either:
1. In expectation over **others'** values (*ex-interim*)
2. In expectation over **all** values (*ex-ante*)

Studied extensively
[Kothari et al., 2003, Archer et al., 2004, Conitzer and Sandholm, 2007, Dekel et al., 2010, Lubin and Parkes, 2012, Mennle and Seuken, 2014, Dütting et al., 2015, Dütting et al., 2017, Azevedo and Budish, 2018, Feng et al., 2018, Golowich et al., 2018]

Literature on γ-IC assumes distribution is known in advance

Where does this knowledge come from?

We relax this assumption:
Assume only **samples** from type distribution

Estimating approximate IC

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

Our estimate:
Maximum utility agent can gain by misreporting her type, on average over samples, when true & reported types from **finite subset** of type space

Estimate used in mechanism design via deep learning:
Add constraint requiring this estimate be small
[Dütting et al., '17, Feng et al., '18, Golowich et al., '18]

Challenge:
Might miss true & reported types with large utility gains

Crucial questions:
1. Which finite subset?
2. What's the estimation error?

Which finite subset?

1. **Uniform grid: Focus of this poster**
   - Easy to construct
   - Works if distribution is "nice"
2. **Learning theoretic cover** (standard from ML theory)
   - Can be hard to construct
   - Always works

Uniform grid: Main challenge

Utility functions are volatile

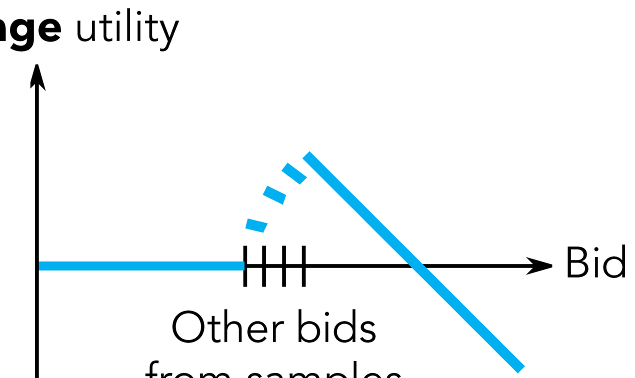

Coarse discretization can lead to poor utility estimates

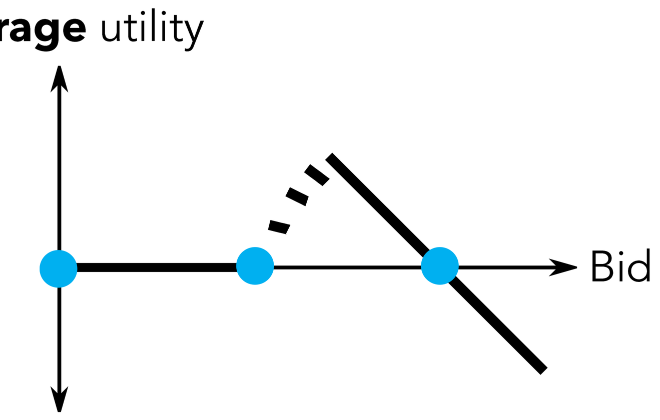

Uniform grid: Guarantees

When is the distribution "nice" enough to use a grid?

Dispersion [Balcan, Dick, and Vitercik, '18]**:**
Functions $u_1, \ldots, u_N$ are *(w,k)-dispersed* if:
Every w-ball contains discontinuities of $\leq k$ functions

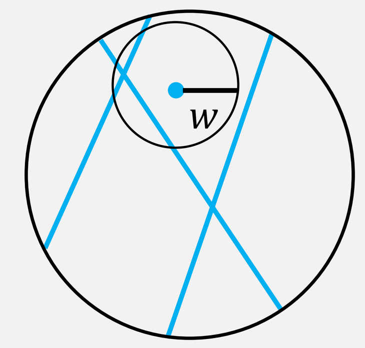

$(w, 2)$-dispersed

Plot average $\frac{1}{N}\sum u_i$:

Not dispersed **Dispersed**

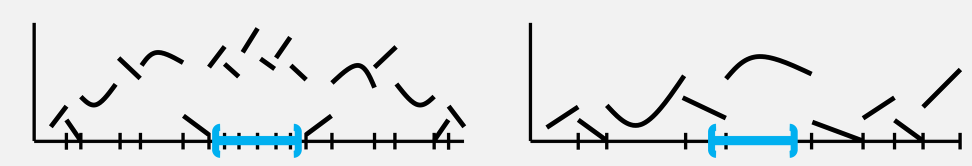

Theorem (informal):
If utility functions induced by N samples are:
1. (w,k)-dispersed
2. Piecewise L-Lipschitz

→ Can use w-grid as finite subset

Estimation error: $\tilde{O}\left(Lw + \frac{k}{N} + \sqrt{\frac{d}{N}}\right)$

d = standard ML measure of utility functions' complexity

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

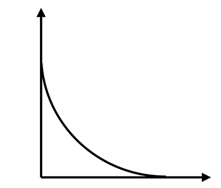

We prove these (w,k) values hold when distribution is **nice**

Applications

$[0, \kappa]$ = range of type distribution's density function

First-price auction
Estimation error = $\tilde{O}\left(\frac{(\#\text{bidders}) + \kappa^{-1}}{\sqrt{(\#\text{samples})}}\right)$

Also analyze **combinatorial** first-price auctions

Generalized second-price auction
Estimation error = $\tilde{O}\left(\frac{(\#\text{bidders})^{3/2} + \kappa^{-1}}{\sqrt{(\#\text{samples})}}\right)$

Discriminatory and uniform price auctions
Generalization of first-price auction to multi-unit settings
Estimation error = $\tilde{O}\left(\frac{(\#\text{bidders})(\#\text{units})^2 + \kappa^{-1}}{\sqrt{(\#\text{samples})}}\right)$