Nonparametric inference on counterfactuals in first-price auctions

P. Andreyanov (HSE) and G. Franguridi (USC)

EEA-ESEM and EARIE August 2022

Introduction

Introduction

A nonparametric estimator of Guerre-Perrigne-Vuong (GPV) and Li-Perrigne-Vuong is typically used for counterfactual analysis

$$\hat{v} = b + \frac{1}{M-1}\frac{\hat{F}(b)}{\hat{f}(b)}, \quad \text{where } F, f \text{ are the cdf, pdf of bids}$$

and M is the number of bidders. Only private values, of course.

Pros:

- does not require (latent) value or cost data
- does not require a parametric specification
- calculates profit margins per-bid
- conceptually simple (hats over f's)
- very fast

Due to separation of profit margins from the covariates, it is only $O(n^2) + O(nc^2)$ compared to, say, NLLS $O(n^2c^2)$, where c is the number of features.

Introduction

Cons:

- only pointwise inference (no uniform)
- density in the denominator - unstable
- still quite slow - roughly $O(n^2)$ without binning

And if I want to have different sizes of auctions,

$$\hat{v} = b + \frac{\sum_{k=1}^{M} p_k \hat{F}^{k-1}(b)}{\sum_{k=1}^{M} p_k (k-1) \hat{F}^{k-2}(b) \hat{f}(b)},$$

where p_k is a share of auctions with k bidders.

Introduction

But what if we use quantiles instead of bids?

$$v(u) = Q(u) + \left[\frac{\sum_{k=1}^{M} p_k u^{k-1}}{\sum_{k=1}^{M} p_k (k-1) u^{k-2}} \right] q(u),$$

where Q, q are the quantile function and quantile density.

Plug-in estimator

$$\hat{v}(u) = \hat{Q}(u) + \left[\frac{\sum_{k=1}^{M} \hat{p}_k u^{k-1}}{\sum_{k=1}^{M} \hat{p}_k (k-1) u^{k-2}} \right] \hat{q}(u).$$

Introduction

$$\hat{v}(u) = \hat{Q}(u) + \left[\frac{\sum_{k=1}^{M} \hat{p}_k u^{k-1}}{\sum_{k=1}^{M} \hat{p}_k (k-1) u^{k-2}} \right] \hat{q}(u)$$

Pros:

- admits varying number of bidders
- no density in the denominator
- still conceptually simple (hats over q's)
- linear in Q, q, and, in fact, just q

But more than that, typical functionals are linear in $v(u)$ and thus, in the quantile density of bids $q(u)$.

Introduction

Typical functionals

$$TS(u^*) = \int_{r^*} x dF^M(x) = \int_{u^*} v(u) du^M$$

$$BS(u^*) = \int_{r^*} (1 - F(x)) F^{M-1}(x) dx = \int_{u^*} (1 - u) u^{M-1} dv(u)$$

$$REV(u^*) = TS(u^*) - M \cdot BS(u^*)$$

Uniform inference on $q(\cdot) \Rightarrow$ uniform inference on TS, BS, REV, where uniform is over the counterfactual exclusion levels u^* (i.e. reserve prices).

This would allow us to test optimality of mechanisms, which in the most natural exercise in auction literature.

Our contribution

Our contribution

A novel procedure for counterfactual analysis

- estimate $\{p_k\}_{k=1}^M$, $Q(u)$ and $q(u)$
- estimate $v(u)$ - the quantile function of latent values
- test uniform hypothesis about functionals that are linear in $v(u)$

We get a substantial improvement over GPV.

- a novel algorithm (two actually) for inference
- the estimator becomes almost pivotal
- elegantly handle random bidders
- lower complexity from $O(n^2)$ to $O(n \log n)$ due to FFT

We prove validity of inference and verify it numerically.
Also, we developed a Python package with a natural fit-predict interface.

Three ingredients

First ingredient

A kernel estimator of quantile density is a weighted average of bid spacings, introduced by Jones in the 80's.

$$\hat{q}_h(u) = \int_0^1 K_h(u-z)\, d\hat{Q}(z) = \frac{1}{n}\sum_{i=1}^{n-1} K_h\left(u - \frac{i}{n}\right)\left(b_{(i+1)} - b_{(i)}\right)$$

- $b_{(i)}$ is the i-th order statistic of $b_1, \ldots, b_n$
- $h > 0$ is bandwidth
- K is a kernel function, $K_h(z) := h^{-1}K(h^{-1}z)$

First ingredient

The computational complexity of this estimator is remarkably small.

$$\hat{q}_h(u) = \int_0^1 K_h(u-z)\, d\hat{Q}(z) = \frac{1}{n}\sum_{i=1}^{n-1} K_h\left(u - \frac{i}{n}\right)\left(b_{(i+1)} - b_{(i)}\right)$$

- sort the bids is $O(n \log n)$
- compute spacings is $O(n)$
- convolve spacings with a static filter is less than $O(n \log n)$

In Python, it is only three lines of code: np.sort, np.diff, np.convolve.
You should not be writing the loop by hand.

Second ingredient

An approximation of the quantile density process by a sample average (cf. classical Bahadur-Kiefer representation) from the 60's

$$\sqrt{nh}\frac{\hat{q}_h(u) - q(u)}{q(u)} = Z_n^*(u) + O_{a.s.}\left(h^{1/2} + h^{-1/2}n^{-1/4}l(n)\right),$$

$$Z_n^*(u) = -\frac{1}{\sqrt{nh}}\sum_{i=1}^{n}\left[K\left(\frac{u - F(b_i)}{h}\right) - \mathbb{E}K\left(\frac{u - F(b_i)}{h}\right)\right],$$

$$Z_n^*(u) \approx -\sqrt{nh}\left[\hat{f}_h^U(u) - \mathbb{E}\hat{f}_h^U(u)\right], \quad \mathbb{E}\hat{f}_h^U(u) \approx 1.$$

where

- $l(n) = (\log n)^{1/2}(\log\log n)^{1/4}$
- $f_h^U(u)$ is the classical KDE estimator with uniform data $F(b_i)$

Second ingredient

Bahadur-Kiefer expansion says

$$\begin{aligned}\frac{\hat{q}_h(u)}{q(u)} - 1 &\approx 1 - \hat{f}_h^U(u) \\ \hat{q}_h^U(u) - 1 &\approx \frac{\hat{q}_h(u)}{q(u)} - 1\end{aligned}$$

To simulate from right-hand side, we just simulate a standard kernel density estimator $\hat{f}_h^U(u)$ with uniform data.

To simulate from left-hand side, we simulate our kernel quantile-density estimator $\hat{q}_h^U(u)$ with uniform data (in fact, with any data) and scale it.

Third ingredient

Until recently, we did not know much about local empirical processes, see Rio (1994), Chernozhukov, Chetverikov & Kato (2014, 2016), to argue the validity of the simulation from either left- and right-handside of BK.

Basically $\hat{q}_h(u), \hat{f}_h^U(u), \hat{q}_h^U(u)$ all diverge (as functions), which rules out direct asymptotic inference via a known limiting distribution.

But, they stay close to each other, if we use the exact same sequence of kernels. Which is not obvious.

This allows for indirect inference via a known sequence of distributions.

Why should you care?

A powerful estimator

Combining the technology from different eras (60's, 80's and 00's) we put forward an estimator that has a powerful combo

- admits pivotal representation
- replaces $O(n^2)$ with $O(n \log n)$
- allows for inference on many natural counterfactuals

A signal to the community

For 20 years GPV has reigned over empirical auctions literature, yet inference on even basic functionals was out of reach. Thus natural hypotheses about optimal mechanisms could not be tested

Pointwise inference on f (the density of valuations) was within reach, but it does not really help.

A signal to the community

It is widely (in my experience) believed that you do not fully estimate the primitives unless you get you hands on the density of valuations f.

In our empirical application, we never even touched the value density, yet, we achieved all of our goals (tested optimality).

Perhaps, $v(u)$ is a better container of model primitives than $f(b)$.

A cool know-how

You wish to test whether a positive reserve price, if set optimally, could increase the revenue.

You plot the revenue as function of the reserve price (actually exclusion levels, or quantiles), and plot uniform confidence bands around it.

If the lower band never gets over the default revenue, this means that the revenue gains at the optimal mechanism are statistically insignificant.

A cool know-how

Here are confidence intervals and bands for the counterfactual revenue from the empirical application with Phil Haile's USFS data.

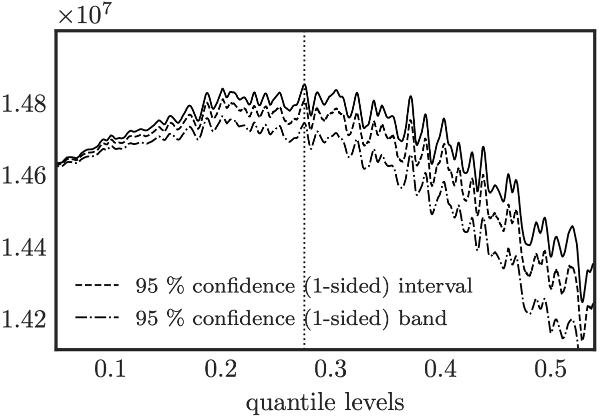

Testing based on confidence intervals is only valid if the researcher is picking the counterfactual reserve price beforehand.

Literature

Literature

- Spacings are cool
  - Ingraham 2005
  - Loertscher, S. and L. M. Marx, 2020
- Highlighted the quantile approach
  - Enache, A. and J.P. Florens, 2012
  - Luo Wan, F, 2018
- Chose "the hard way" with the f density
  - Marmer, V et al 2012, 2019, 2021
  - Zincenko, F, 2021

Thank you!