Optimal Robust Double Auctions (dynamic but with private values)

Pasha Andreyanov (HSE)
Junrok Park (Taiwan University)
Tomasz Sadzik (UCLA)

EWMES (Berlin), Dec 2022

Double auctions

* means robust

	single-unit	multi-unit
one-sided	FPA, English, SPA*, Dutch*	VCG*, GSP, Ausubel*
double	McAffee*	uniform-price, pay-as-you-bid, Incentive, ???

- uniform-price has 0 budget
- McAffee and Incentive almost break even
- VCG in 2-sided setting will actually produce deficit
- UP, PAYB and Incentive are all non-robust

This paper

In an environment with many buyers and sellers that have private information and multi-unit demand, but independent utilities...

...we develop a optimal (rev. maximizing) robust (ex-post IC, IR) mechanism and implement it in a dynamic fashion.

Goal 1: find optimal direct mechanism

Goal 2: find double-clock implementation

Goal 3: eliminate unwanted equilibria

Goal 4: some quant. results

Let me unpack this...

Part 1: find optimal robust direct mechanism

- ex post revelation principle, virtualization
- locate the (ex-post) worst-off type
- prove a small lemma[1]
- recast the problem as convex optimization

Part 2: find (double-)clock implementation

- merge Ausube with double-clock design
- assess the stop-off price
- introduce marginal tax $m\tau$
- redefine sincere bidding
- establish sincere equilibrium

[1] defeat non-monotonicity of virtual type at the (ex-post) worst-off type

Part 3: minimize unwanted equilibria

- understand incentives to deviate
- define informational spillover
- prove a small lemma [1]
- optimize over price paths

Part 4: get some quant. results

- pick a tractable specification (quadratic)
- pick some distribution (uniform, logistic)
- compare to progressive tax (σ-vcg)[2]
- compare to bid-ask spread (δ-vcg)[2]

[1] either there is spillover only on one side or there is spillover for at most one agent
[2] see Andreyanov Sadzik (2021) for σ-vcg and δ-vcg mechanisms.

Several points

Point 1: without robustness non-trivial shading

Thus turn to robust

Point 2: two-sided private information + efficient = budget deficit

Thus turn to optimal

Point 3: two-sided $\Rightarrow$ non-monotone virtual type

Some kind of ironing will happen at worst-off-type

Point 4: robust + 2-sided = wot endogenous

Eventually, this will be fixed

Point 5: path-dependent payments + dynamic $\Rightarrow$ bad equilibria

Thus optimize over price paths

Part 1: find optimal robust direct mechanism

Consider a classical single-crossing (but non-linear) consumption utility $u(\theta, q)$ and quasi-linear payoff $u(\theta, q) - t(\theta)$ structure. Let the type θ and allocation q be 1-dim.

Leading example: $u(\theta, q) = \theta q - \mu q^2$, where μ is known

What is an optimal (revenue-maximizing) robust mechanism?

- denote (equilibrium) surplus as $s(\theta) = u(\theta, q(\theta)) - t(q(\theta), \theta_{-i})$
- maximize average $u(\theta, q) - s(\theta)$
- ex post IC constraint $s(\theta) = \max_q[u(\theta, q) - t(q, \theta_{-i})]$
- ex post IR constraint $s(\theta) \geqslant u(\theta_i, 0)$

Standard virtualization techniques apply

$$v_i(\theta_i, q|\theta_{-i}) = u(\theta_i, q) - \frac{I(\theta_i > wot(\theta_{-i})) - F(\theta_i)}{f(\theta_i)} \frac{\partial}{\partial \theta_i}[u(\theta_i, q) - u(\theta_i, 0)]$$

where $\text{wot}(\theta_{-i})$ is the ex post worst-off type.

Leading example: $v(\theta_i, q|\theta_{-i}) = (\theta - \frac{I(\theta_i > wot(\theta_{-i})) - F(\theta_i)}{f(\theta_i)})q - \mu q^2$.

Problems:

- this is not even a "private utility"
- virtual utility is (downwards) discontinuous in own type
- worst-off type is endogenous to the sought mechanism

Seems we have a fixed-point type of problem with ironing on top...

... but don't panic!

A useful lemma: let $\text{tet}(\theta_{-i})$ be the type excluded from trade, that is, who trades exactly zero, then it is one of the worst-off types $\text{wot}(\theta_{-i})$.

in other words, $tet(\theta_{-i}) \subset wot(\theta_{-i})$.

Proof: the slope of surplus is linked to the sign of your trade by single crossing, thus if sub-gradient contains zero it is also the argmin.

Thus virtual utility can be rewritten wlog

$$v(\theta_i, q) = u(\theta_i, q) - \frac{I(q > 0) - F(\theta_i)}{f(\theta_i)} \frac{\partial}{\partial \theta_i} [u(\theta_i, q) - u(\theta_i, 0)]$$

Which is continuous in type (at the cost of kink at $q = 0$) so no ironing needed. Most importantly - endogeneity is solved, it is a "private utility".

All we have to do is maximize sum of virtual utilities, or, equivalently, find and enforce an "efficient" allocation that of the virtual economy.

Part 2: find double-clock implementation

How to design the auction?

- Robustness implies Vickrey-style transfers
- Vickrey + Dynamic = Ausubel (clinching) design
- presumably, two Ausubel auctions running towards each other while continuously clearing the market, until the clock prices meet

But how about the budget deficit?

Small detour to Andreyanov Sadzik (2021) paper.

Lets add "quadratic" tax for each transaction, to subsidize the loss associated with Vickrey-style payments.

- marginal tax $m\tau(q) = \sigma q$, where $\sigma > 0$.
- integrated tax $\tau(q) = \sigma q^2/2$

Bidding is sincere in the sense that bidders play truthfully, but as if their utility was deformed $u(q) \to u(q) - \tau(q)$, and the auction finds the efficient allocation in the deformed economy.

Alternatively add a "bid-ask spread"

- marginal tax $m\tau(q) = \delta sign(q)$, where $\delta > 0$.
- integrated tax $\tau(q) = \delta|q|$

But what is optimal $m\tau$?

We need to somehow make the bidder perceive her utility as $v(\theta_i, q)$ instead of $u(\theta_i, q)$, so that in equilibrium she behaves as if she maximizes the sum of virtual utilities.

- Denote $m\tau(p, q) := mu_i(\hat{\theta}, q) - mv_i(\hat{\theta}, q)$
- And $\hat{\theta}(p, q)$ is such that $p = mv_i(\hat{\theta}, q)$, trust me

At the conjectured allocation, the first order condition is always met:

$$p = mv_i(\theta_i, q) = mu(\theta_i, q) - m\tau(p, q).$$

Second order conditions are shown "by inspection".

But additional assumptions will be required to tame $v(\theta_i, q)$.

Namely, log-concave distro (like logistic or uniform) and control the signs of $u'''_{qq\theta}$ and $u'''_{q\theta\theta}$ (the latter hold trivially for the leading example).

Back to the implementation

So let's run two Ausubel auctions: forward and reverse, with clocks moving towards each other, with $m\tau(p, q)$ on top.

Then there will be a sincere equilibrium in the twisted sense.

But if we found $m\tau(p, q)$ correctly it will maximize the sum of virtual utilities and the stop-off price is the WE price in the virtual economy.

Demands revealed are:

$$d(p) = argmax_q[u(\theta_i, q) - \tau(p, q) - pq].$$

Transfers made are:

$$t(q) = \int_0^q p_{-i}(x) + m\tau(p_{-i}(x), x)dx$$

where $p_{-i}(q)$ is the residual curve facing agent i.

Zoom in closer on $m\tau$

The optimal (revenue maximizing) tax has to

- depend on the clock price $\tau(p, q)$
- likely regressive (concave) rather than progressive (convex)
- integrated tax $\tau(p, q)$ exhibit a kink at $q = 0$
- equivalently, $m\tau(p, q)$ exhibit a discontinuity at $q = 0$
- does not depend on the number of bidders

Here is an illustration

Let the utility be quadratic and private type distributed $U[-1, 1]$.

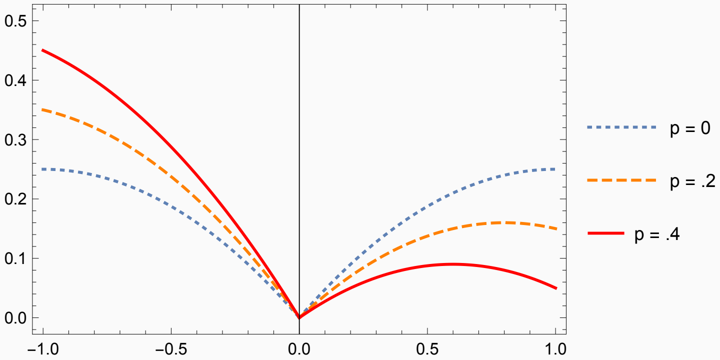

Plot of $\tau(q, p)$ with q on the horizontal axis.

Two key features

- kink at zero creates exclusion of weak traders
- shoulders minimize distortion for strong traders

You can also think of it in terms of the marginal tax, instead of integrated tax, but it is less intuitive geometrically.

Again the utility be quadratic and private type distributed $U[-1, 1]$.

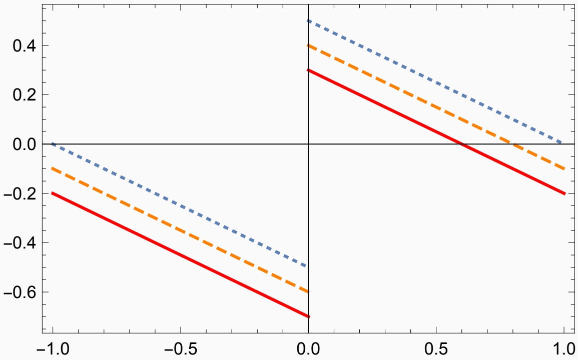

Plot of $m\tau(q, p)$ with q on the horizontal axis, for different prices.

But who computes the mechanism?

- Designer only computes $m\tau(p, q)$ from the knowledge of u, F, and two non-linear equations, separately for each bidder
- Player only solves foc: $mu(q) - m\tau(p, q) = p$
- The mechanism (and all the endogenous worst-off-types) are computed by the equilibrium outcomes of the auction like in some human circuit-board.

Exclusion pattern

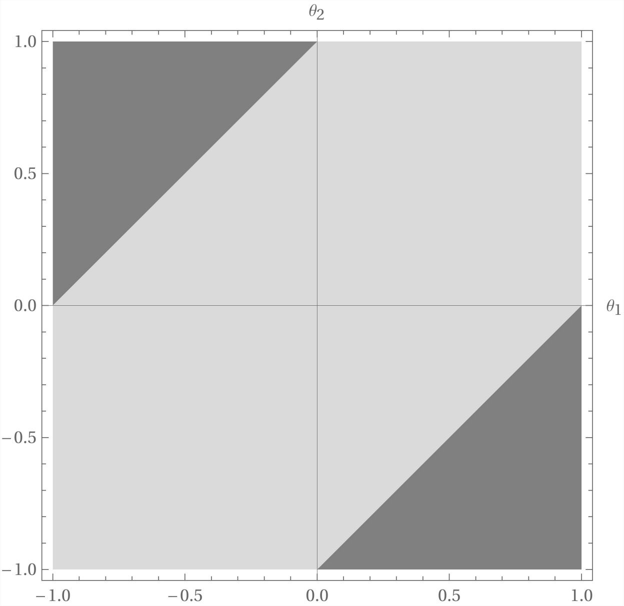

Axes: θ_1, θ_2. Gray - exclusion, dark gray - no exclusion.

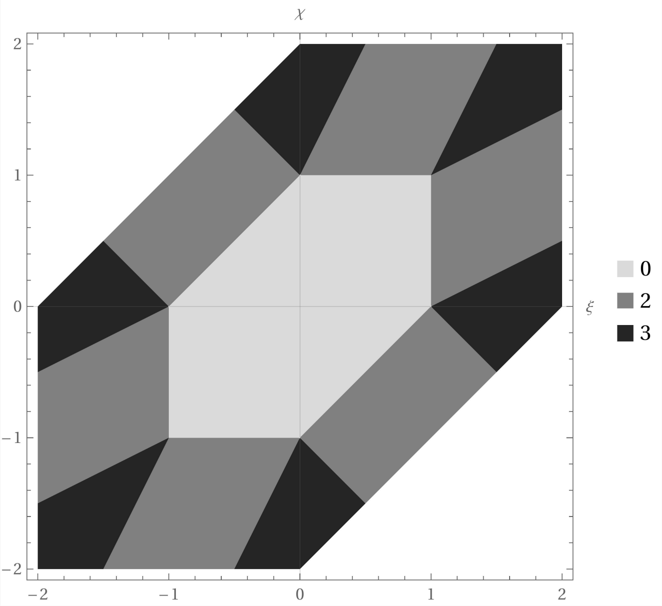

Axes: $\xi = \theta_1 - \theta_3$, $\chi = \theta_2 - \theta_3$. Gray - full exclusion, dark gray - partial exclusion, black - no exclusion.

Part 3: eliminate unwanted equilibria

Fact: Existence of sincere equilibrium does not depend on the price path. But existence of other equilibria does.

Imagine that you fully advance the forward auction till the stop-off price (when markets clear), and then you run the reverse.

- All players in the reverse auction are aware of the stop-off price, as well as their final allocation now.
- So they can immediately drop their demands to that allocation.

Advancing demands prematurely makes the transfers less favorable for the auctioneer. This may create strong incentives to collude.

Another example

- I play in the forward auction and my sincere demand is $q_i^{fwd}(p) > 0$
- At the same time in the reverse auction I clinch something.
- Once I am a confirmed seller, I know that I will never buy.
- Replace my sincere demand with $q_i^{fwd}(p) = 0$ in the forward auction.
- My payoffs do not change
- But payoffs of others will change (vcg style)

Thus there exist certain strategies, that are even lazier than sincere bidding, which undermines our confidence in the desired equilibrium.

Let me summarize

- bidders may learn the stop-off price ahead of time
- bidders may learn if they are buyers/sellers ahead of time

This is a consequence of a more general phenomenon - the inadvertent informational spillover between the forward and reverse auctions.

Let $p^+ \leqslant p^-$ be clock prices in forward and reverse auctions. Let q_i^+ be the revealed demand in forward and $q_i^-(p)$ in the reverse auction.

Let $q_{-i}^+ = -\sum_{j \neq i} q_j^+$ the residual demand in forward auction, and $q_{-i}^- = -\sum_{j \neq i} q_j^-$ the residual demand in reverse auction.

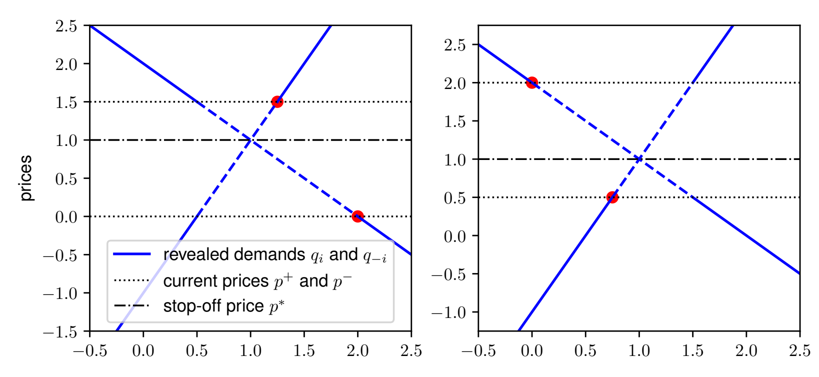

Left figure - spillover into forward auction. Right figure - spillover into reverse auction.

Let's define it formally

- Agent i experiences spillover into forward auction iff $q_i^+ > q_{-i}^-$
- Agent i experiences spillover into reverse auction iff $q_i^- < q_{-i}^+$

Can we move the clocks to minimize spillovers?

It is a pure engineering problem, no game theory here.

Good news: if we prevent spillover entirely, this will lead to uniqueness (via iterated elimination) of sincere equilibrium, eventually.

Bad news: with 2 players, preventing spillover entirely is not possible.

What is possible then?

A useful lemma: either there is spillover into only one auction (forward or reverse) or there is spillover for at most one agent.

Put differently, you can minimize the total number of spillovers by balancing the spillovers into forward and reverse auctions, eventually having only 1 or 0 agents exposed.

We want to reach and maintain this ecosystem for as long as possible.

Proof: assume that there is spillover into both auctions, and also for different agents

- $-\sum_{k\neq i} q_k^+ = q_{-i}^+ > q_i^-$ for some i
- $q_j^+ > q_{-j}^- = -\sum_{k\neq j} q_k^-$ for some $j \neq i$

Then

$$-\sum_{k\neq i,j} q_k^+ \geqslant q_i^- + q_j^+ \geqslant -\sum_{k\neq i,j} q_k^-$$

or

$$\sum_{k\neq i,j} q_k^+ < \sum_{k\neq i,j} q_k^-$$

which contradicts monotonicity of demand ($q_k^- \leqslant q_k^+$).

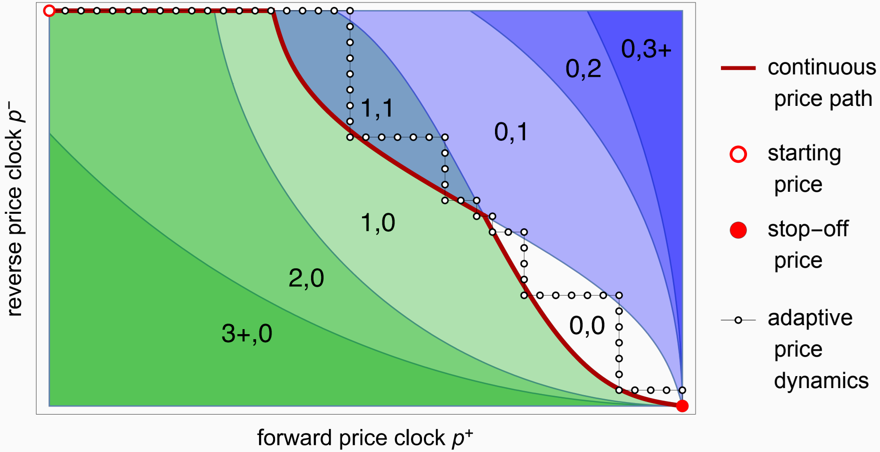

Adaptive price dynamics: if there are more spillovers into forward - move forward clock, if there are more spillovers into reverse - move reverse clock, if equal move either clock.

Part 4: some quant. results

To maximize tractability assume

$$u(\theta, q) = \theta q - \frac{q^2}{2}$$

Pick one of two distributions: uniform or logistic

Pick three families of mechanisms:

- pareto frontier $J_\alpha(\theta, q) = \alpha u(\theta, q) + (1 - \alpha) v(\theta, q)$
- smooth taxation $J_\sigma(\theta, q) = u(\theta, q) - \sigma q^2/2$, call it σ-vcg
- pure exclusion $J_\delta(\theta, q) = u(\theta, q) - \delta |q|$, call it δ-vcg

Plot the three frontiers

With uniform distribution, smooth taxation extracts approx 50% of optimal revenue, while pure exclusion approx 75%.

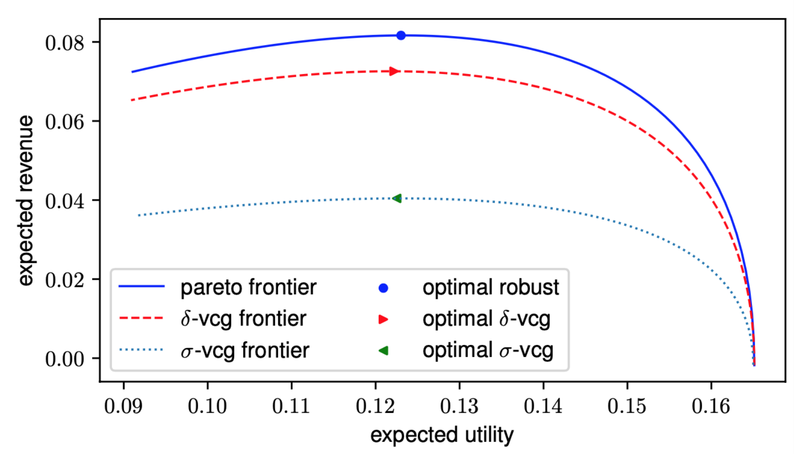

for all target levels of efficiency, apparently

With logistic distribution, smooth taxation extracts approx 65% of optimal revenue, while pure exclusion approx 99%.

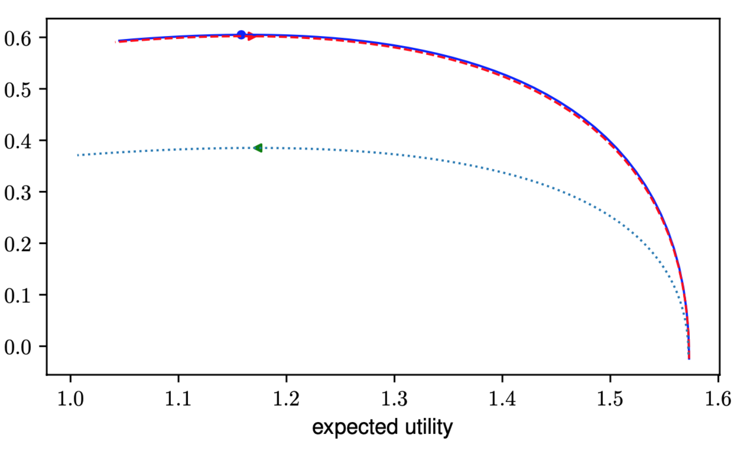

for all target levels of efficiency, apparently

Conclusion

- direct mechanism is difficult to solve for
- but not because of ironing (tet $\subset$ wot)
- indirect mechanism is easy to solve for, surprisingly
- double-clock design + vickrey style payments allow for optimization over price paths so to minimize informational spillovers
- if $m\tau(p, q)$ is too much, use pure exclusion instead