Going... going... wrong: a test of level-k (and cognitive hierarchy) models of bidding behaviour

Itzhak Rasooly
University of Oxford

Preliminaries

- Accompanying website: `https://auctionsolver.herokuapp.com`
- A working paper is also online

Introduction

In recent years, the level-k auction model (Crawford and Iriberri, 2007) has emerged as a rival to more traditional equilibrium based approaches:

- Level-0 types either bid their valuation or bid randomly
- For $k \geq 1$, level-k bids optimally on the assumption that all opponents are level $k - 1$
- Most individuals assumed to be levels $1 - 3$

Introduction

Lots of debate over whether level-k can explain the 'winner's curse' in common value auctions:

- Crawford and Iriberri (2007)
- Ivanov et al. (2010)
- Costa-Gomes and Shimoji (2015)

Introduction

Growing interest in the implications of the level-k auction model for mechanism design:

- Crawford et al. (2009)
- De Clippel et al. (2019)

Introduction

Surprisingly, however, there is little evidence on whether level-k can outperform equilibrium in the basic IPV setting that forms the starting point for most auctions research:

- Most experiments use uniform values. But then level-k and equilibrium exactly coincide (for all $k \geq 1$)

Introduction

Surprisingly, however, there is little evidence on whether level-k can outperform equilibrium in the basic IPV setting that forms the starting point for most auctions research:

- Most experiments use uniform values. But then level-k and equilibrium exactly coincide (for all $k \geq 1$)
- Crawford and Iriberri (2007)'s IPV data do not strongly separate the models

Introduction

Surprisingly, however, there is little evidence on whether level-k can outperform equilibrium in the basic IPV setting that forms the starting point for most auctions research:

- Most experiments use uniform values. But then level-k and equilibrium exactly coincide (for all $k \geq 1$)
- Crawford and Iriberri (2007)'s IPV data do not strongly separate the models
- Kirchkamp and Reiß (2010)'s setting better separates the models (marginally), but their analysis of level-k is incorrect

Introduction

Surprisingly, however, there is little evidence on whether level-k can outperform equilibrium in the basic IPV setting that forms the starting point for most auctions research:

- Most experiments use uniform values. But then level-k and equilibrium exactly coincide (for all $k \geq 1$)
- Crawford and Iriberri (2007)'s IPV data do not strongly separate the models
- Kirchkamp and Reiß (2010)'s setting better separates the models (marginally), but their analysis of level-k is incorrect
- Non-experimental approaches (Gillen, 2009; An, 2017) cannot be used to convincingly decide between the models

Introduction

Plan for paper:

1. Find *simple* environments that dramatically disentangle the two models' predictions
2. Construct these environments in the laboratory and check which model performs the best

Main (substantive) finding: despite its success in other domains, the level-k model cannot explain behaviour in auctions

The all-pay auction

Set-up:

- $n \geq 2$ bidders
- Everybody pays their bid (even the losers)
- Values and bids are restricted to $\mathbb{X} = \{0, 1, 2, ..., x\}$ where $x \in \mathbb{N}^+$
- Nobody wins if the highest bid is a tie (as in the experiment)
- For the level-k analysis, will assume that values are uniformly and independently distributed

The all-pay auction: equilibrium analysis

Proposition 1

The discrete all-pay auction has exactly one symmetric Bayes-Nash equilibrium.

Sketch of proof.

To prove existence, just apply Harsanyi (1967). To prove uniqueness, argue that any SE must be *monotone*, *gapless* and satisfy $P(b = 0|v = 0) = 1$. Given these facts, any candidate SE strategy can be equivalently represented in *jump form*. One can then define an inductive algorithm which produces the candidate SE; and finally show that no other strategy could be an SE. □

The all-pay auction: equilibrium analysis

- Importantly, the proof gives us recipe for computing the SE: `https://auctionsolver.herokuapp.com/`
- In the case of uniform values, we already know that the equilibrium is in properly mixed strategies (Rasooly and Gavidia-Calderon, 2020)
- However, the algorithm tells us that expected bids are well-approximated by the continuous formula

$$\beta(v) = \left(\frac{n-1}{n}\right)\frac{v^n}{(x-1)^{n-1}}$$

The all-pay auction: equilibrium analysis

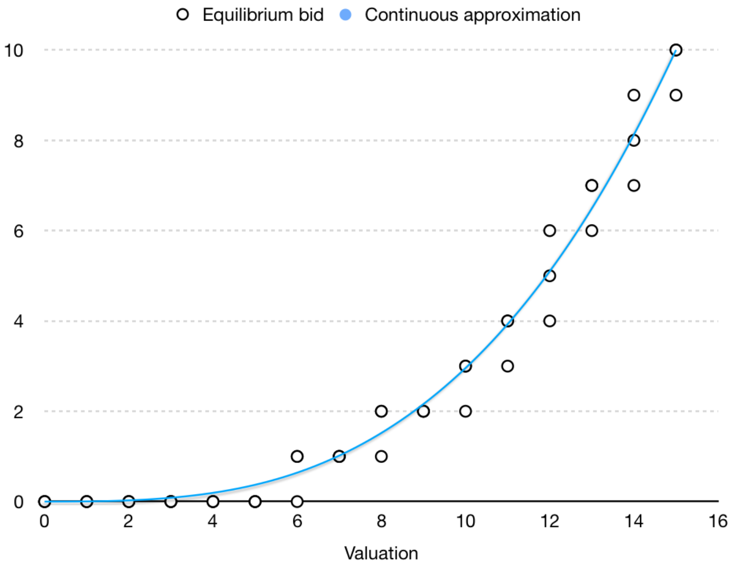

The all-pay auction: level-k analysis

Given that level-0 bids are uniform, the level-1 player solves the problem

$$\max_{b \in \mathbb{X}} \quad \pi(v, b) \equiv v\mathrm{P}(\text{win}|b) - b = v\left(\frac{b}{x+1}\right)^{n-1} - b.$$

The all-pay auction: level-k analysis

Given that level-0 bids are uniform, the level-1 player solves the problem

$$\max_{b \in \mathbb{X}} \quad \pi(v, b) \equiv v\mathrm{P}(\text{win}|b) - b = v\left(\frac{b}{x+1}\right)^{n-1} - b.$$

To solve this, suppose first that $n = 2$. Then

$$\pi(v, b) \equiv b \underbrace{\frac{v}{x+1}}_{<1} - b$$

So $b^* = 0$ (for all v!)
If this is true for $n = 2$, it must also be true for all $n \geq 2$.

The all-pay auction: level-k analysis

- As we have seen, $\beta^1(v) = 0$ for all $v \in \mathbb{X}$

The all-pay auction: level-k analysis

- As we have seen, $\beta^1(v) = 0$ for all $v \in \mathbb{X}$
- So $\beta^2(v) = 1$ for $v \geq 2$ and $\beta^2(v) = 0$ otherwise

The all-pay auction: level-k analysis

- As we have seen, $\beta^1(v) = 0$ for all $v \in \mathbb{X}$
- So $\beta^2(v) = 1$ for $v \geq 2$ and $\beta^2(v) = 0$ otherwise
- Similarly, $\beta^3(v) = 2$ for $v \geq 3$ and $\beta^3(v) = 0$ otherwise

The all-pay auction: level-k analysis

- As we have seen, $\beta^1(v) = 0$ for all $v \in \mathbb{X}$
- So $\beta^2(v) = 1$ for $v \geq 2$ and $\beta^2(v) = 0$ otherwise
- Similarly, $\beta^3(v) = 2$ for $v \geq 3$ and $\beta^3(v) = 0$ otherwise
- etc. etc.
- Note: we are assuming that level-k players break ties by choosing the lowest optimal bid.

The all-pay auction: level-k analysis

What happens at very high levels?

- Given our assumption about tie-breaking, level-k can never coincide with the symmetric equilibrium

The all-pay auction: level-k analysis

What happens at very high levels?

- Given our assumption about tie-breaking, level-k can never coincide with the symmetric equilibrium
- Moreover, level-k must cycle as $k \to \infty$
  - If it did not cycle, there would need to exist some $k \in \mathbb{N}$ such that β_K coincides with β'_K for all $k' > k$
  - But then β_K would be a symmetric, pure strategy equilibrium, contradicting Rasooly and Gavidia-Calderon (2020)'s Prop 4.

The all-pay auction: level-k analysis

Proposition 2

Suppose that

$$k \leq (x+1)^{\frac{n-1}{n}} + 1.$$

Then a level-k bidder sets $\beta^k(v) = k - 1$ *for* $v \geq k$ *and* $\beta^k(v) = 0$ *otherwise. Moreover, there is no* $k \in \mathbb{N}$ *at which* β^k *coincides with equilibrium; with the implication that* β^k *cycles as* $k \to \infty$.

The all-pay auction: comparing the models

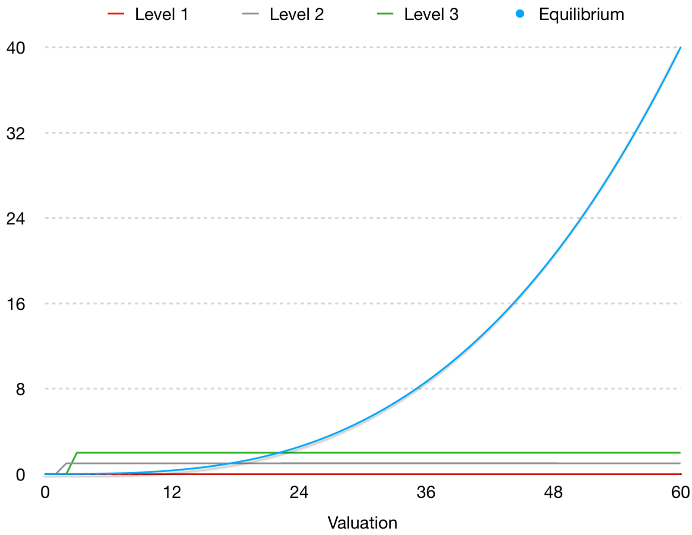

The first-price auction

Set-up:

- $n \geq 2$ bidders
- **Only the winner pays their bid**
- Values and bids are restricted to $\mathbb{X} = \{0, 1, 2, ..., x\}$ where $x \in \mathbb{N}^+$
- Nobody wins if the highest bid is a tie
- Values are uniformly and independently distributed
- **Each player is only allowed to bid with probability p**

The first-price auction: equilibrium analysis

Proposition 3

The discrete first-price auction has exactly one symmetric Bayes-Nash equilibrium if $p \in (0, 1)$.

Sketch of proof.

As before! □

The first-price auction: equilibrium analysis

- As before, the proof tells us exactly *how* to compute the (unique) SE
- With uniform values, expected bids remain well-approximated by the continuous equilibrium, which is now

$$\beta(v) = \left(\frac{n-1}{n}\right) v - \frac{x(1-p)}{np}\left[1 - \left(\frac{1-p}{1-p+p(v/x)}\right)^{n-1}\right]$$

The first-price auction: equilibrium analysis

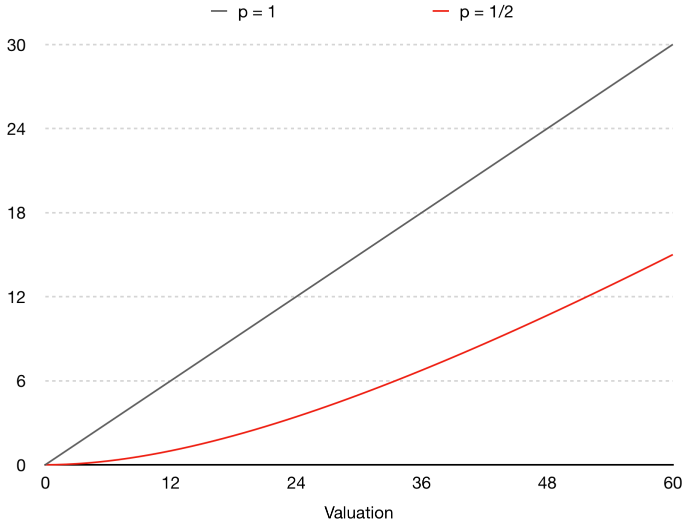

The first-price auction: level-k analysis

Given that level-0 bids are either cancelled or uniform on $\mathbb{X}$, a level-1 player solves the problem

$$\max_{b \in \mathbb{X}} \quad (v - b)\mathrm{P}(\mathrm{win}|b) = (v - b)\left(1 - p + \frac{pb}{x+1}\right)^{n-1}$$

Routine optimisation then reveals that

$$\beta^1(v) \approx \max\left\{\left(\frac{n-1}{n}\right)v - \left(\frac{1-p}{p}\right)\frac{x}{n},\ 0\right\}$$

The first-price auction: level-k analysis

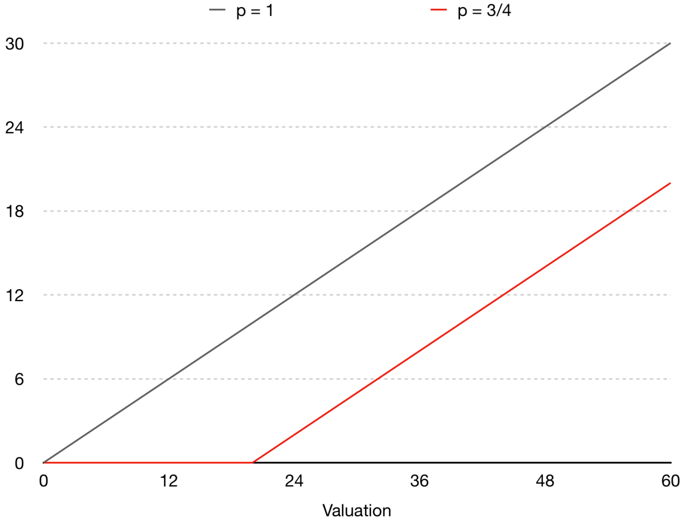

The first-price auction: separating the models

Our goal is to disentangle the two theories. To this end, we define the distance between the theories as

$$d(p) \equiv \int_0^x \left|\beta(v) - \beta^1(v)\right| dv.$$

We now study the probability p^* that maximises this distance.

The first-price auction: comparing the models

Proposition 4

$p^* \in [1/n, 1)$.

The first-price auction: separating the models

In fact, a numerical analysis reveals that $p^* \approx 1/n$.

Table: Optimal cancellation probabilities (rounded)

n	2	3	4	5	6	7	8
p^*	0.536	0.343	0.256	0.204	0.170	0.145	0.127

The first-price auction: comparing the models

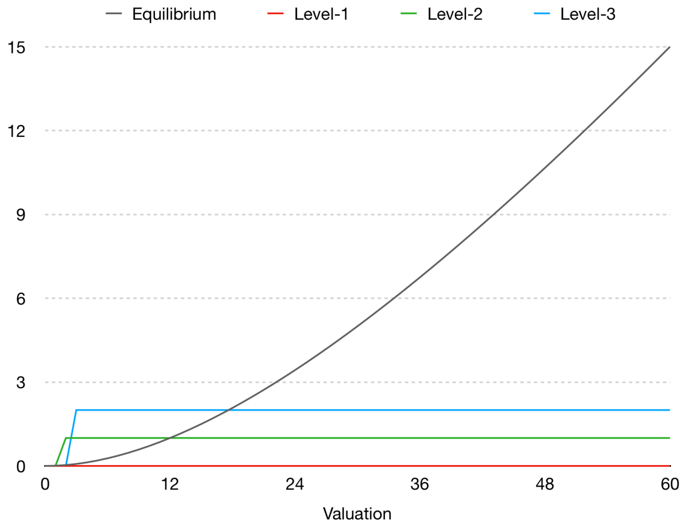

The first-price auction: level-k cycling

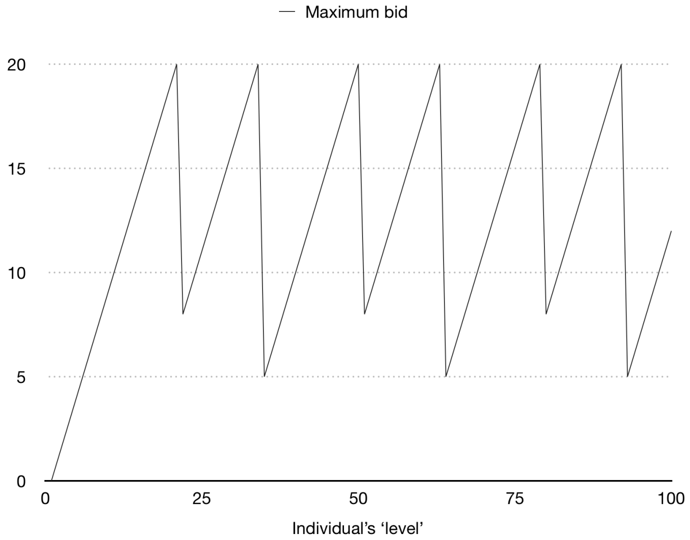

Experimental design

Main idea:

1. Create the aforementioned auction structures in the 'lab' (using induced values)
2. Check which theory (if either) can explain observed bidding

Experimental design

Some details:

- Experiment written in oTree and conducted online
- Subjects recruited by CESS
- Incentivised (some auctions randomly selected to 'count') and pre-registered

Experimental design

More details:

- Each subject played 2 rounds of every auction format
- They were not informed of the bids in previous rounds
- ‘Perfect stranger matching’
- Subjects needed to pass extensive quizzes before they could proceed to either auction
- The experiment concluded by calibrating levels using Alaoui and Penta (2016)’s variant on Arad and Rubinstein (2012)’s 11/20 game and measuring risk aversion using the BRET

Experimental results: first-price auction

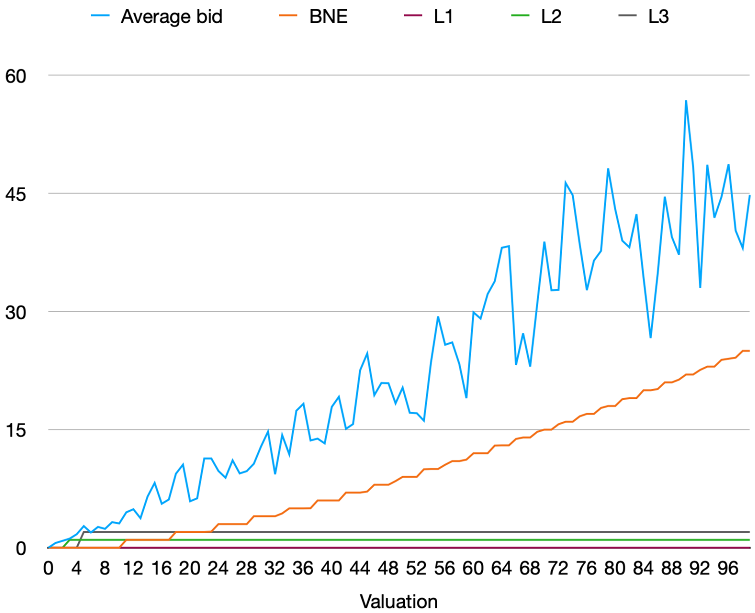

Experimental results: all-pay auction

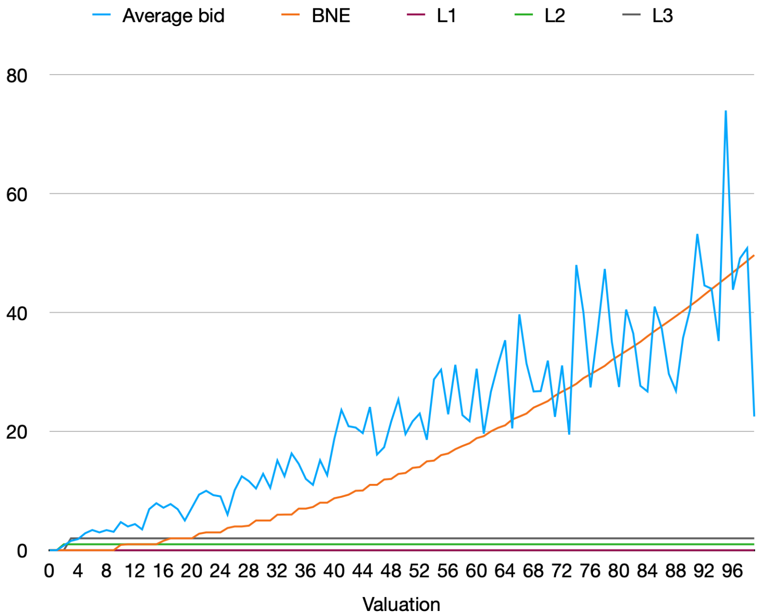

Experimental results: error rates

Table: Prediction errors

	T1 FP	T1 AP	T2 FP	T2 AP
BNE	15.1	13.3	14.1	14.3
Level-k	21.3	19.8	21.4	24.0

Experimental results: likelihoods

Table: Comparing equilibrium and level-k

		T1FP	T1AP	T2FP	T2AP
BNE	LL	-4407.9	-4332.7	-1024.7	-1042.2
	BIC	8820.5	8670.2	2053.1	2088.1
L1	LL	-4530.4	-4501.2	-1088.8	-1128.6
	BIC	9065.4	9007.1	2181.3	2260.9
L1-L2	LL	-4530.4	-4501.2	-1085	-1126.9
	BIC	9070.1	9011.7	2177.4	2261.2
L1-L3	LL	-4530.1	-4501.2	-1071.1	-1120.7
	BIC	9074.3	9016.4	2153.4	2252.6

Conclusions

- One can design simple (so implementable) environments that dramatically disentangle level-k and equilibrium predictions
- In these environments, level-k predictions + comparative statics are highly implausible
- Why does the model fail so badly while performing so well in other settings? Perhaps (1) lack of salient anchor (2) computational complexity

Conclusions

- *Pace* Crawford et al. (2009) and De Clippel et al. (2019), the level-k model should not be used for auction design
- While equilibrium can be hard to believe in, ‘behavioural’ models can be even worse!

Alaoui, L. and Penta, A. (2016). Endogenous depth of reasoning. The Review of Economic Studies, 83(4):1297–1333.

An, Y. (2017). Identification of first-price auctions with non-equilibrium beliefs: A measurement error approach. Journal of econometrics, 200(2):326–343.

Arad, A. and Rubinstein, A. (2012). The 11-20 money request game: A level-k reasoning study. American Economic Review, 102(7):3561–73.

Costa-Gomes, M. A. and Shimoji, M. (2015). A comment on "can relaxation of beliefs rationalize the winner's curse?: An experimental study". Econometrica, 83(1):375–383.

Crawford, V. P. and Iriberri, N. (2007). Level-k auctions: Can a nonequilibrium model of strategic thinking explain the winner's curse and overbidding in private-value auctions? Econometrica, 75(6):1721–1770.

Crawford, V. P., Kugler, T., Neeman, Z., and Pauzner, A. (2009). Behaviorally optimal auction design: Examples and observations. Journal of the European Economic Association, 7(2-3):377–387.

De Clippel, G., Saran, R., and Serrano, R. (2019). Level-k mechanism design. The Review of Economic Studies, 86(3):1207–1227.

Gillen, B. (2009). Identification and estimation of level-k auctions. Available at SSRN 1337843.

Harsanyi, J. C. (1967). Games with incomplete information played by "bayesian" players, i–iii part i. the basic model. Management Science, 14(3):159–182.

Ivanov, A., Levin, D., and Niederle, M. (2010). Can relaxation of beliefs rationalize the winner's curse?: An experimental study. Econometrica, 78(4):1435–1452.

Rasooly, I. and Gavidia-Calderon, C. (2020). The importance of being discrete: on the (in-) accuracy of continuous approximations in auction theory. arXiv preprint arXiv:2006.03016.