Structural Estimation of Directional Dynamic Games With Multiple Equilibria

Fedor Iskhakov, Australian National University
Dennis Kristensen, University College London
John Rust, Georgetown University
Bertel Schjerning, University of Copenhagen

Department of Econometrics and Business Statistics
Monash University
December 3 2025

Directionality $\implies$ Full solution $\implies$ Nested MLE

- Focus on the subclass of discrete stochastic games with directional state transitions $\to$ directional dynamic games
- State recursion algorithm = generalization of backward induction to solve for MPE in stages
- Multiple solutions at stage games $\to$ RLS algorithm to enumerate all combinations of solutions $\iff$ solve for all MPEs (IRS, 2016)
- This paper develops a nested full solution MLE estimator $\equiv$ NRLS estimator based on integer programming branch-n-bound algorithm
  - Fully robust to multiplicity of equilibria in the model and the data
  - Computationally feasible
  - Computational burden decreases with sample size
- Extensive Monte Carlo study of existing estimators: 2-step, NPL, EPL, MPEC vs NRLS

Iskhakov, Rust and Schjerning (2016, ReStud)
Recursive Lexicographical Search: Finding All Markov Perfect Equilibria of Finite State Directional Dynamic Games

MLE for dynamic games with multiple equilibria

- Data from M independent markets from T periods, N players

$$\mathrm{Z} = \left\{a^{ijt}, x^{it}\right\}_{i\in\{1\ldots M\}, j\in\{1,\ldots,N\}, t\in\{1,\ldots,T\}}$$

- MPE is a pair of strategy profiles and value functions such that

$$\begin{aligned} V_\theta &= \Psi^{V_\theta}(V_\theta, P_\theta, \theta) && \text{(Bellman equations)} \\ P_\theta &= \Psi^{P_\theta}(V_\theta, P_\theta, \theta) && \text{(CCPs = mutual best responces)} \end{aligned}$$

- Multiplicity $\rightarrow$ set of equilibria $\mathcal{E}(\theta) = \{V_\theta^k, P_\theta^k\}_{k\in\{1,\ldots,K(\theta)\}}$
- MLE estimator $\hat{\theta}^{ML}$ is given by

$$\hat{\theta}^{ML} = \arg\max_\theta \left[\max_{k\in\{1,\ldots,K(\theta)\}} \sum_{t=1}^{T} \sum_{i=1}^{M} \sum_{j=1}^{N} \log P_j^k(a^{ijt} | x^{it}; \theta) \right]$$

- One equilibrium in the data, relax later with grouped fixed effects
- Inner loop requires full solution, impossible?

MLE by contrained optimization (MPEC)

- Idea: use discretized values of P and V as *variables*
- Augmented log-likelihood function

$$\mathcal{L}(Z, P, \theta) = \sum_{t=1}^{T} \sum_{i=1}^{M} \sum_{j=1}^{N} \log P_j(a^{ijt} | x^{it}; \theta)$$

- The constrained optimization formulation of the ML estimation problem is

$$\max_{\theta, P, V} \mathcal{L}(Z, P, \theta) \text{ subject to } \begin{cases} V = \Psi^V(V, P, \theta) \\ P = \Psi^P(V, P, \theta) \end{cases}$$

- May work with multiple equilibria with smart optimization algorithms
- Much bigger computational problem
- Implements the same MLE estimator (*when it works*)

Su (2013); Egesdal, Lai and Su (2015)

Other existing estimation methods

- Two step (CCP) estimators
  - Fast, do not impose equilibrium constraints, finite sample bias
  1. Estimate CCP $\to \hat{P}$
  2. Method of moments • Minimal distance • Pseudo likelihood

  Hotz, Miller (1993); Altug, Miller (1998); Pakes, Ostrovsky, and Berry (2007); Pesendorfer, Schmidt-Dengler (2008)

- Nested pseudo-likelihood (NPL)
  - Recursive two step pseudo-likelihood
  - Bridges the gap between efficiency and tractability
  - Unstable under multiplicity

  Aguirregabiria, Mira (2007); Aguirregabiria, Marcoux (2021)

- Efficient pseudo-likelihood (EPL)
  - Incorporates Newton step in the NPL operator
  - More robust to the stability and multiplicity of equilibria

  Dearing, Blevins (2024)

Directional dynamic games (DDG)

DDG is a finite state stochastic game where state transitions under all feasible Markovian strategies form a directional acyclic graph with self-loops
(see IRS, 2016 for formal definition)

- points of finite state space $X \to$ vertexes of the graph
- Edge from x_i to $x_j \iff$ under some strategy profile the hitting probability of x_j from x_i is positive

1. Simple algorithm to determine if graph is a DAG
2. Topological sort to find totally ordered partition
3. Backward recursion on the found total order

DDGs in the literature:

Judd, Schmedders, Yeltekin (2012); Dube, Hitsch, Chintagunta (2010); Iskhakov, Rust, Schjerning (2018); Anderson, Rosen, Rust, Wong (2024)

Graph representation of all possible transitions on X

Transitions induced by all feasible strategy profiles

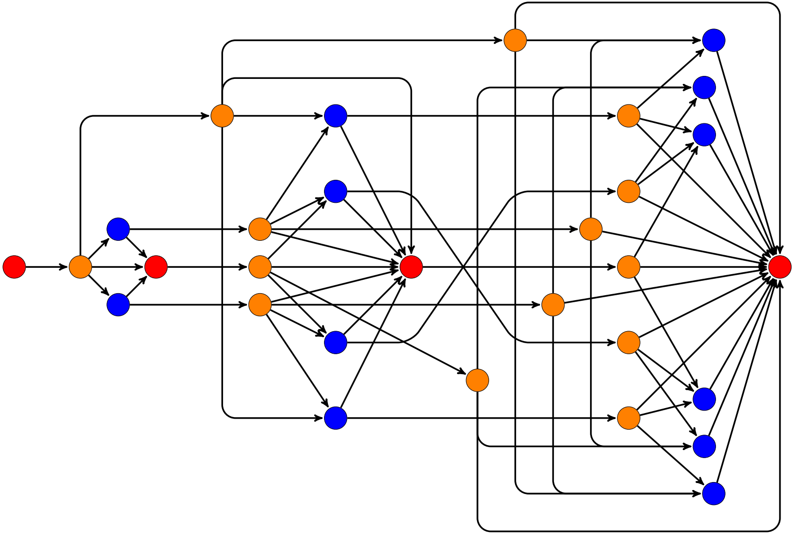

Total order on the partition of the state space

After running a topoligical sort algorithm on the DAG

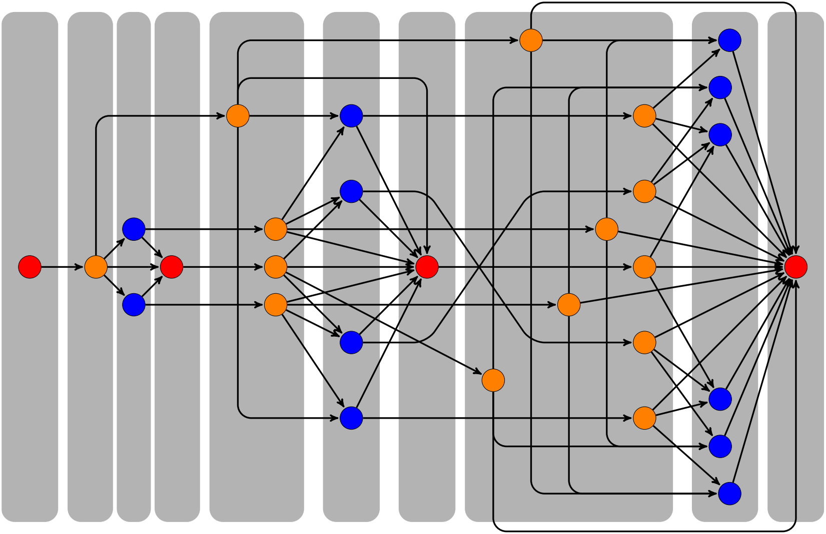

State recursion = backward induction on the state space

Solving subgames in continuation strategies → many small problems

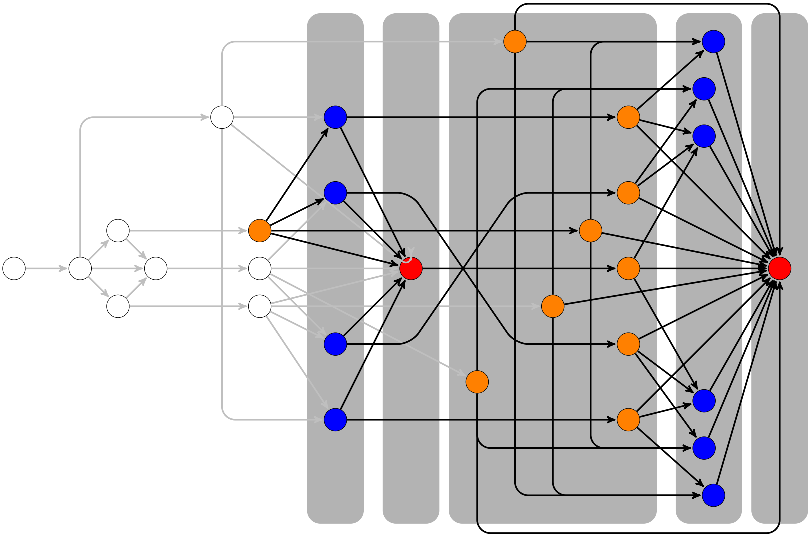

Multiplicity of stage equibiria $\Longleftrightarrow$ Multiplicity of MPE

- State recursion proceeds conditional on equilibrium selection rule
- Selected equilibrium at downstream stage affects the equilibria and number of equilibria at upstream stages
- Need to systematically combine different stage equilibria $\rightarrow$ Recursive Lexicographical Search $\equiv$ Depth-first tree traversal

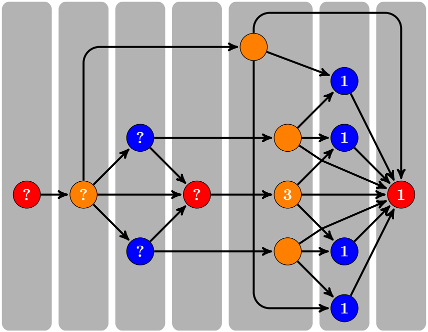

RLS = tree traversal

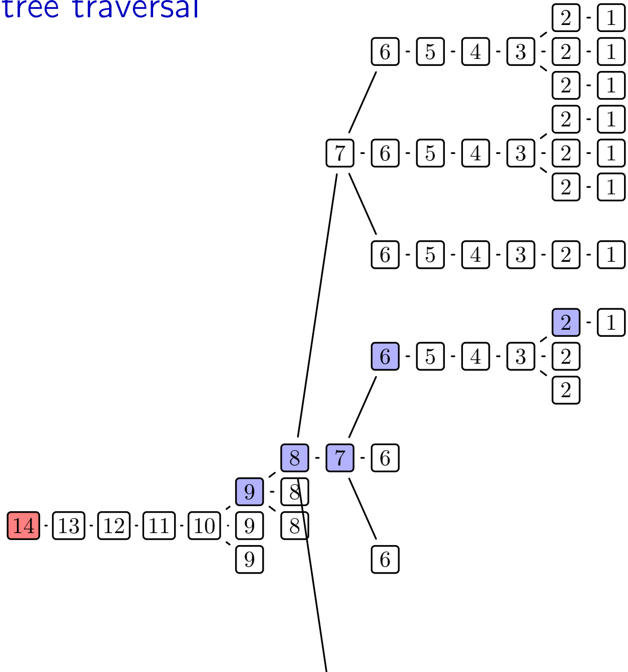

Nested Recursive Lexicographical Search (NRLS)

- Data from M independent markets from T periods, N players

$$Z = \left\{a^{ijt}, x^{it}\right\}_{i\in\{1\dots M\}, j\in\{1,\dots,N\}, t\in\{1,\dots,T\}}$$

- Set of equilibria $\mathcal{E}(\theta) = \{V_\theta^k, P_\theta^k\}_{k\in\{1,\dots,K(\theta)\}}$

1. Outer loop Maximization of the likelihood function w.r.t. to structural parameter θ

$$\theta^{ML} = \arg\max_{\theta\in\Theta} \mathcal{L}(Z,\theta)$$

2. Inner loop Maximization of the likelihood function w.r.t. equilibrium selection $\equiv$ discrete parameter $k \in \{1,\dots,K(\theta)\}$

$$\mathcal{L}(Z,\theta) = \arg\max_{k\in\{1,\dots,K(\theta)\}} \mathcal{L}(Z,\theta,P_\theta^k)$$

- With multiple equilibria in the data $\mathcal{L}(Z,\theta)$ has more elaborate form

Likelihood over the state space

Can efficiently represent likelihood by counts of observations

- With equilibrium k choice probabilities $P_j^k(a|x;\theta)$, likelihood is

$$\mathcal{L}(Z,\theta,P_\theta^k) = \sum_{t=1}^{T}\sum_{i=1}^{M}\sum_{j=1}^{N} \log P_j^k(a^{ijt}|x^{it};\theta)$$

- Let ι index points in the state space
  $\iota = 1$ in the initial subset, $\iota = |X|$ in the terminal subset of X
- Denote n_ι the number of observations in state x_ι and $n_\iota^{a_j}$ the number of observations of player i taking action a_j at x_ι

$$n_\iota = \sum_{t=1}^{T}\sum_{i=1}^{M} \mathbb{1}\{x^{it} = x_\iota\} \qquad n_\iota^{a_j} = \sum_{t=1}^{T}\sum_{i=1}^{M} \mathbb{1}\{a^{ijt} = a_j, x^{it} = x_\iota\}$$

- Then equilibrium-specific likelihood is given by

$$\mathcal{L}(Z,\theta,P_\theta^k) = \sum_{\iota=1}^{|X|}\sum_{j=1}^{N}\sum_{a_j} n_\iota^{a_j} \log P_j^k(a_j|x_\iota;\theta)$$

Branch and bound solution method

Land and Doig, 1960 *Econometrica*

- Old method for solving integer programming problems

- **Branching**: RLS tree
- **Bounding**: The bound function is partial likelihood of equilibrium k calculated on the subset of states $\iota \in \mathcal{S} \subset X$

$$\mathcal{L}^{\text{part}}(Z^{\mathcal{S}}, \theta, P^k_\theta) = \sum_{\iota \in \mathcal{S}} \sum_{j=1}^{N} \sum_{a_j} n_\iota^{a_j} \log P_j^k(a_j | x_\iota; \theta)$$

- Where $Z^{\mathcal{S}} = \{(a, x) : x \in \mathcal{S}\}$ denotes data observed on $\mathcal{S}$
- Monotonic decreasing in cardinality of $\mathcal{S}$ (declines as more data is added)
- Equals to the full log-likelihood on the full state space when $Z^{\mathcal{S}} = Z$ (at the leafs of RLS tree, next slide)

BnB on RLS tree, step 1

$\boxed{14} \cdot \boxed{13} \cdot \boxed{12} \cdot \boxed{11} \cdot \boxed{10}$ Partial loglikelihood = -3.2

BnB on RLS tree, step 2

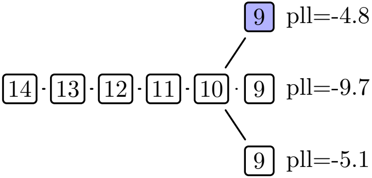

BnB on RLS tree, step 3

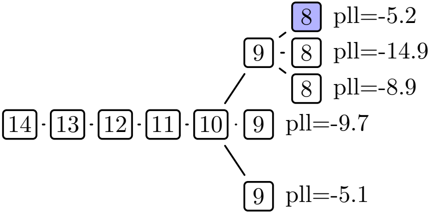

BnB on RLS tree, step 4

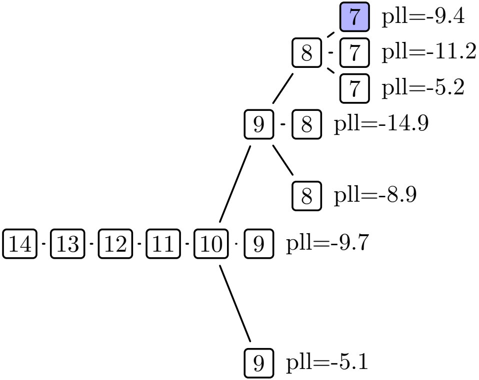

BnB on RLS tree, step 5

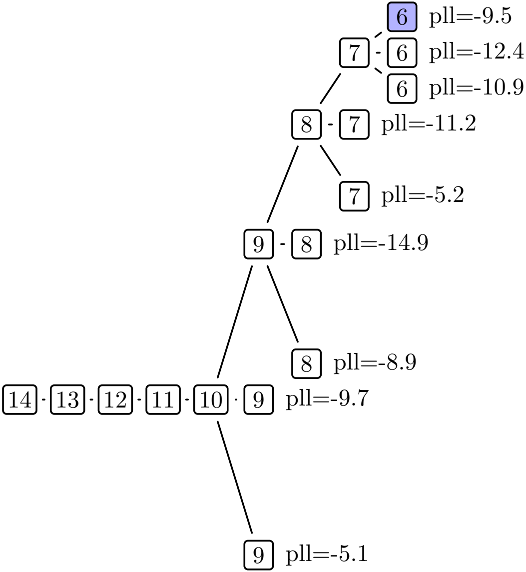

BnB on RLS tree, step 6

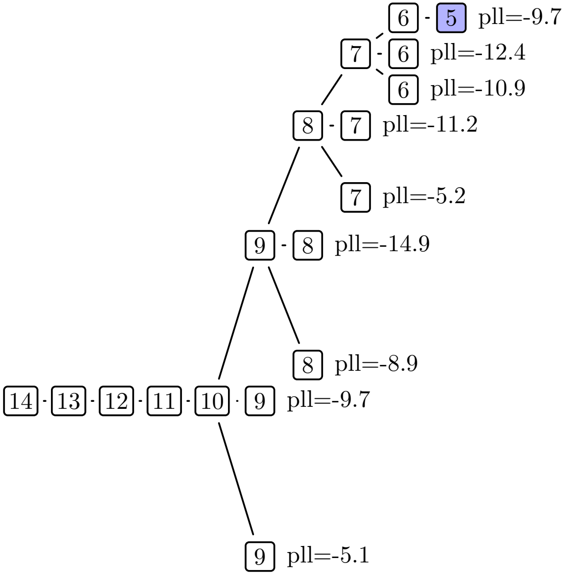

BnB on RLS tree, step 7

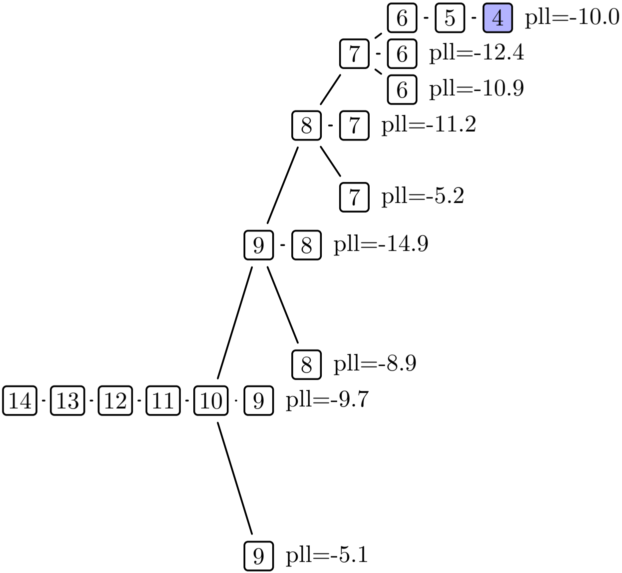

BnB on RLS tree, step 8

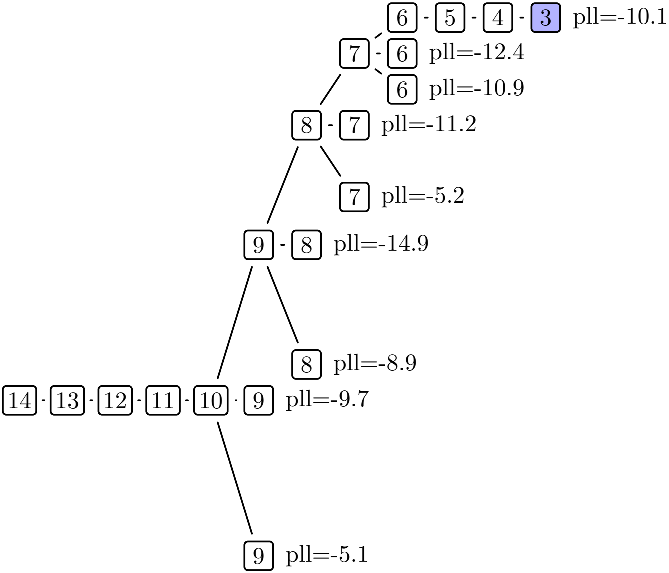

BnB on RLS tree, step 9

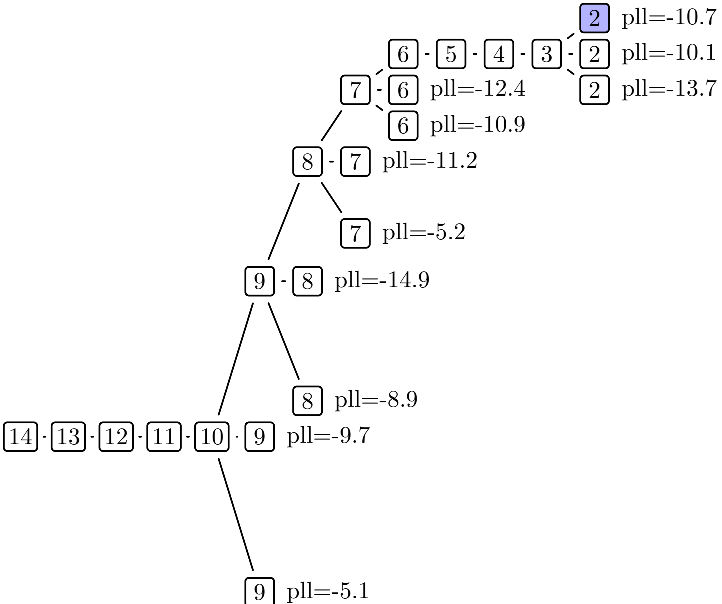

BnB on RLS tree, step 10

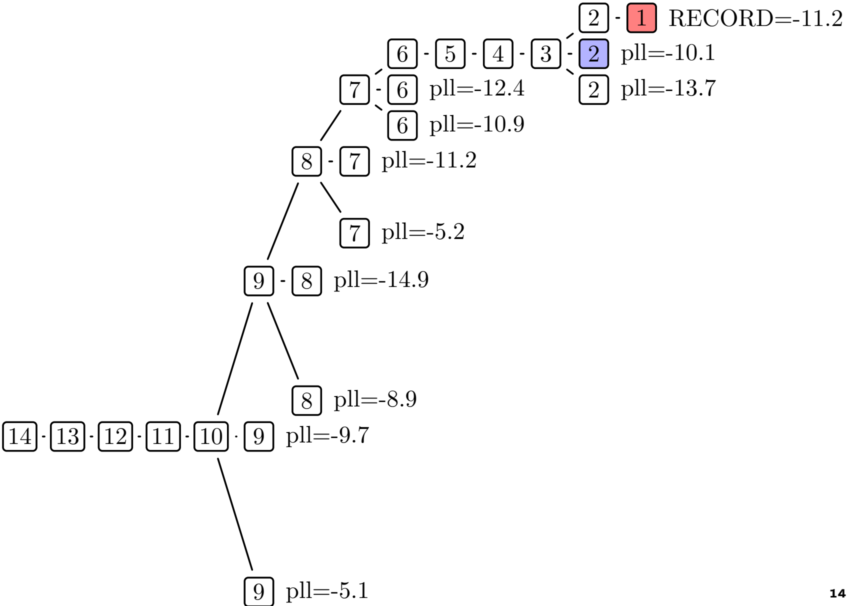

BnB on RLS tree, step 11

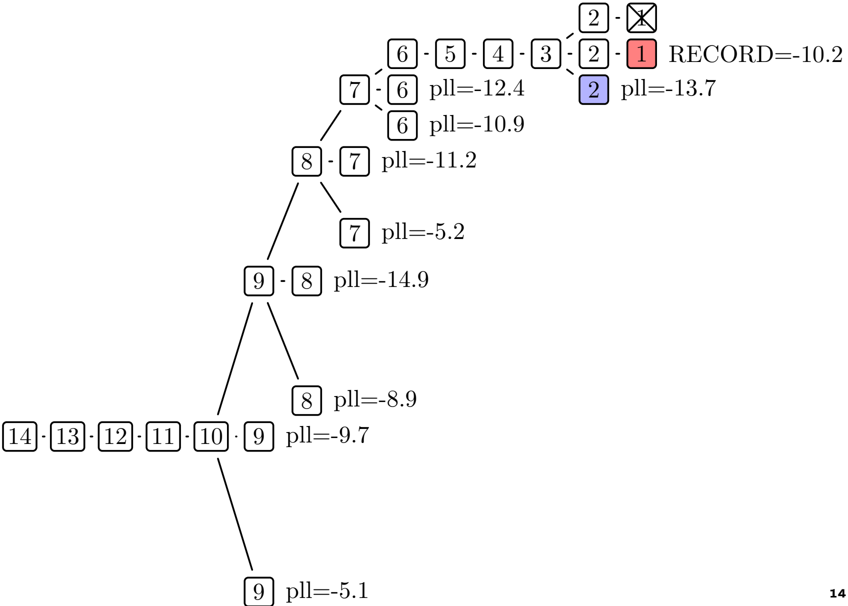

BnB on RLS tree, step 12

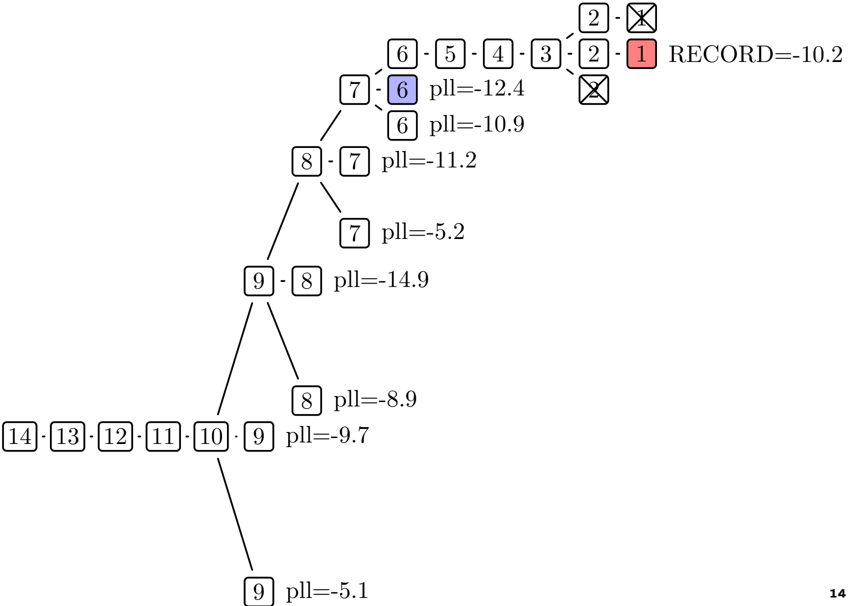

BnB on RLS tree, step 28

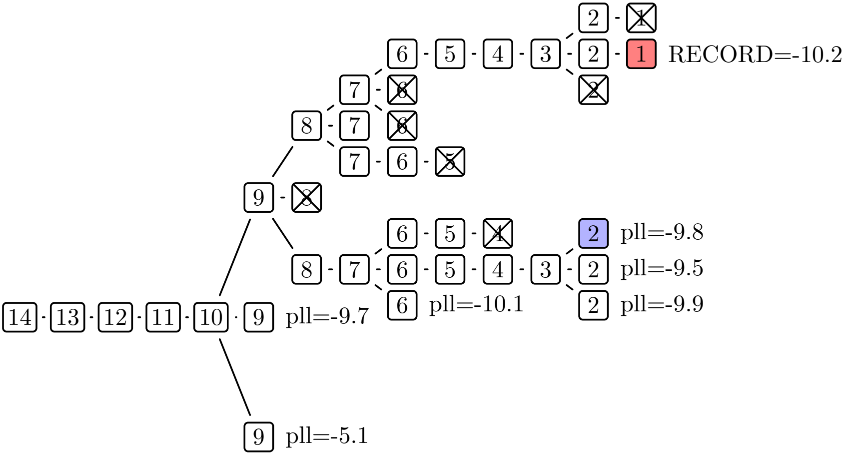

BnB on RLS tree, step 29

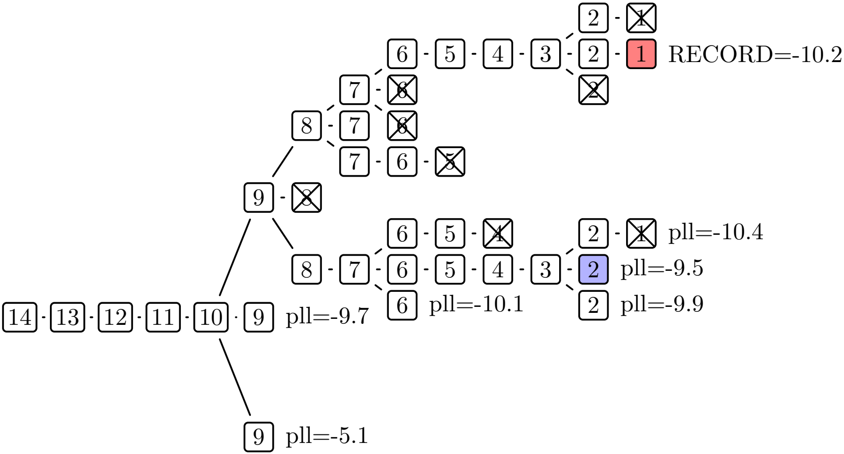

BnB on RLS tree, step 30

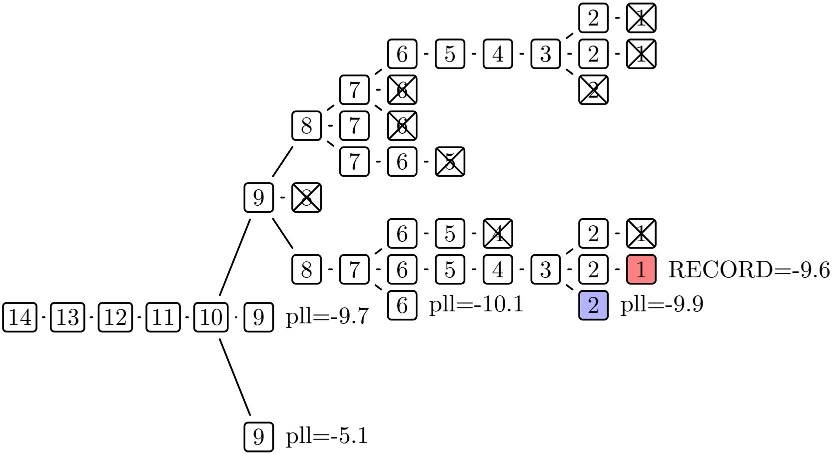

BnB on RLS tree, step 31

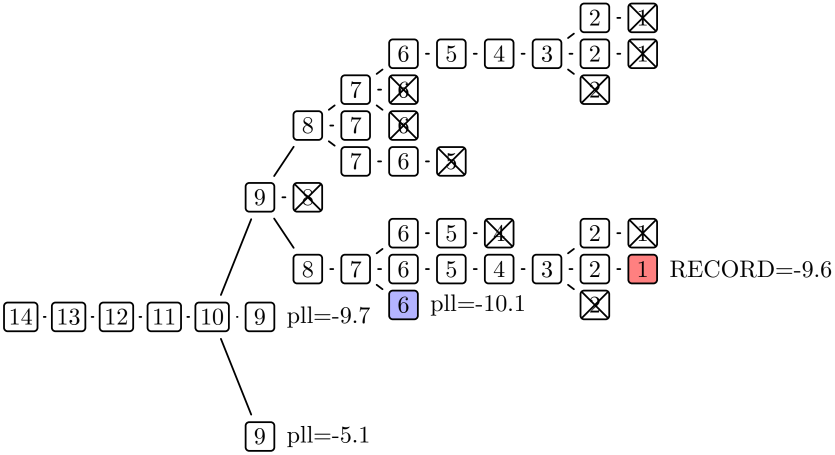

BnB on RLS tree, step 34

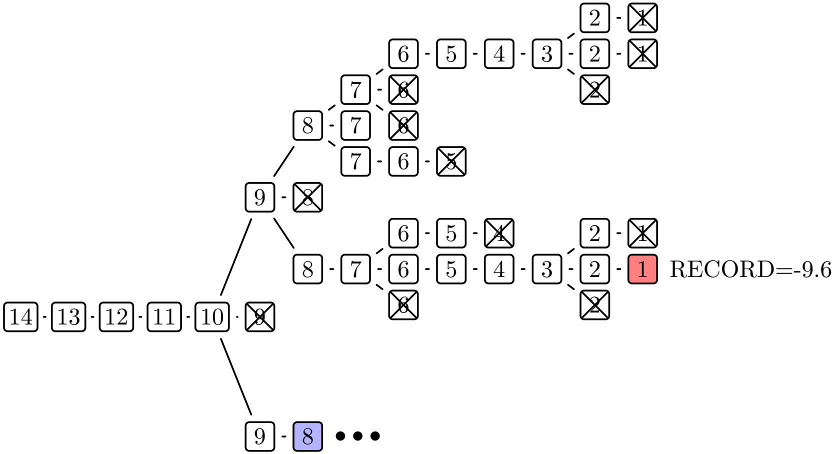

Refinement: non-parametric likelihood bounding

Much more powerful bound for BnB based on empirical frequencies

- Replace choice probabilities $P_j^k(a_j|x_\iota;\theta)$ with frequencies $n_\iota^{a_j}/n_\iota$

$$\mathcal{L}^{\text{non-par}}(Z^{\mathcal{S}}) = \sum_{\iota\in\mathcal{S}} \sum_{i=1}^{J} \sum_{a} n_\iota^{a_i} \log(n_\iota^{a}/n_\iota)$$

- $\mathcal{L}^{\text{non-par}}(Z^{\mathcal{S}})$ depends only on the counts from the data!
- Not hard to show algebraically that for any $Z^{\mathcal{S}}$ (≈Gibbs inequality)

$$\mathcal{L}^{\text{non-par}}(Z^{\mathcal{S}}) > \mathcal{L}^{\text{part}}(Z^{\mathcal{S}}, \theta, P_\theta^k) \; \forall \theta, k$$

- Therefore partial likelihood can be optimistically extrapolated by empirical likelihood at any step ι of the RLS tree traversal

$$\mathcal{L}^{\text{part}}(Z^{\{|X|,\dots,\iota\}}, \theta, P_\theta^k) + \mathcal{L}^{\text{non-par}}(Z^{\{\iota-1,\dots,1\}})$$

Non-parameteric likelihood bounding

$\iota = |X| = 14$ (terminal state) on the left, $\iota = 1$ (initial state) on the right

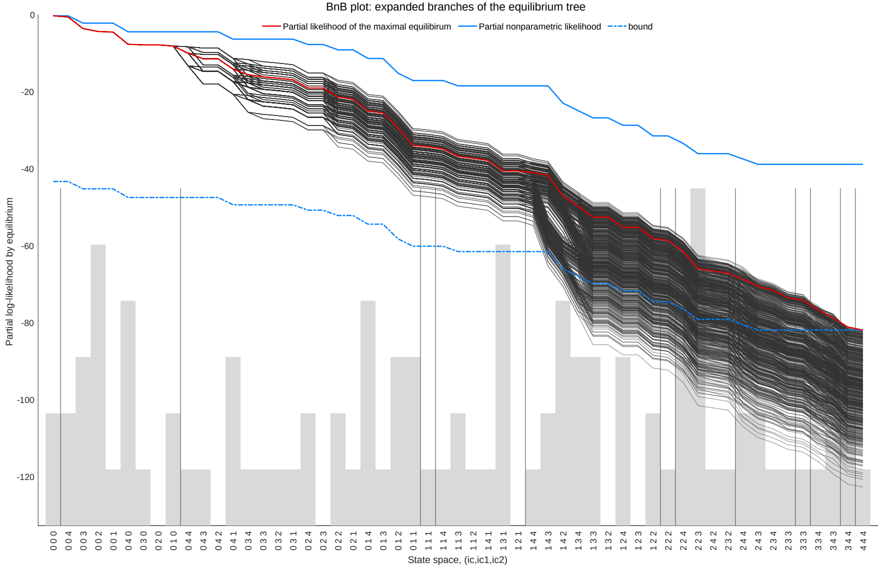

BnB with non-parameteric likelihood bound

Greedy traversal + non-parameteric likelihood bound

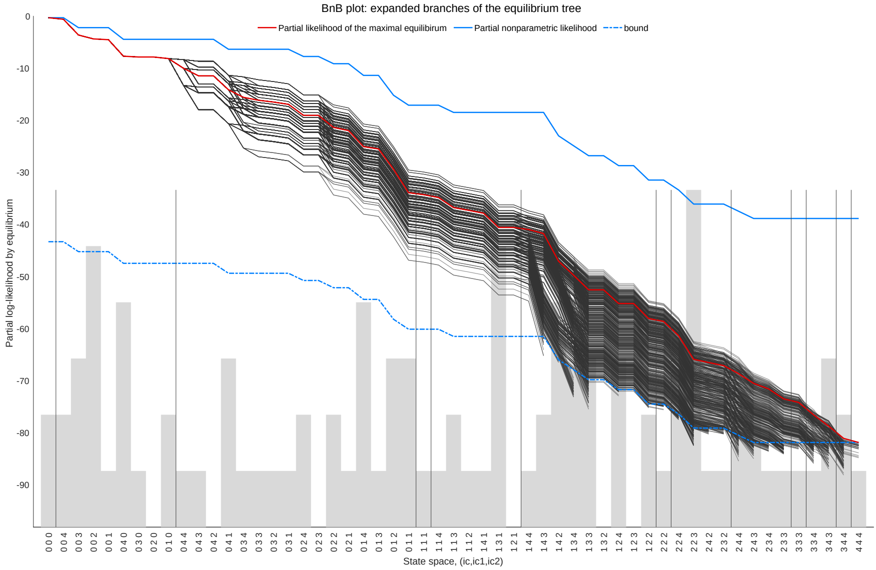

BnB with non-parameteric likelihood bound, larger sample

Non-parametric → parametric likelihood as $N \to \infty$ at true $\theta \Rightarrow$ even less computation

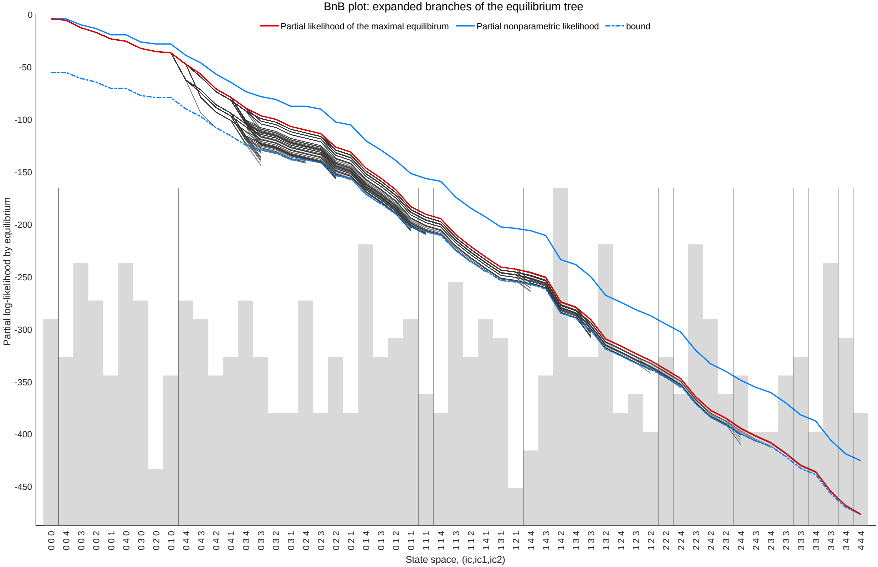

BnB refinement with non-parametric likelihood

- For any amount of data the non-parametric likelihood is greater or equal to the parametric likelihood *algebraically*
- BnB augmented with non-parameteric likelihood bound gives sharper Bounding Rules $\rightarrow$ less computation

- Wih more data as $M \rightarrow \infty$, $T \rightarrow \infty$
- Non-parametric log-likelihood converges to the actual likelihood
- The width of the band between the blue lines in the plots decreases

  $\rightarrow$ Even sharper Bounding Rules
  $\rightarrow$ Even less computation

BnB yields exact solution of the inner integer maximization problem
$\Longrightarrow$ MLE for any sample size, but easier to compute with more data!

Monte Carlo simulations

A

Single equilibrium in the model
One equilibrium in the data

B

Multiple equilibria in the model
Same equilibrium played the data

C

Multiple equilibria in the model
Multiple equilibria in the data:

- Long panels, each market plays their own equilibrium
- Groups of markets play the same equilibrium

Implementation details:

- Leapfrogging model with $N = 2$ Bertrand competitors deciding whether to invest in cost-reducing technology (IRS, 2016)
- k_1 parameter in investment cost function
- $M = 1000,\ T = 5$
- All methods are initialized with 2-step CCP estimator

Monte Carlo A: no multiplicity

Number of equilibria at true parameter: 1
Number of equilibria in the data: 1

	2step	NPL	EPL	MPEC-VP	MPEC-P	NRLS
True $k_1 = 3.5$	3.52786	3.49714	3.49488	3.49488	3.49486	3.49488
Bias	0.02786	-0.00286	-0.00512	-0.00512	-0.00514	-0.00512
MCSD	0.10037	0.06522	0.07042	0.07042	0.07078	0.07042
ave log-like	-1.16661	-1.16144	-1.16143	-1.16143	-1.16139	-1.16143
log-likelihood	-5833.07	-5807.21	-5807.16	-5807.16	-5806.95	-5807.16
log-like short	-	-0.050	-0.000	-0.000	-0.000	-0.000
KL divergence	0.03254	0.00021	0.00024	0.00024	0.00024	0.00024
$\|\|P - P0\|\|$	0.11270	0.00469	0.00495	0.00495	0.00500	0.00495
$\|\|\Psi(P) - P\|\|$	0.16185	0.0000	0.0000	0.0000	0.0000	0.0000
$\|\|\Gamma(v) - v\|\|$	0.87095	0.00000	0.00000	0.00000	0.00000	0.00000
Converged of 100	-	100	100	100	99	100

- Equilibrium conditions satisfied (except 2step)
- Nearly all MLE estimators identical to the last digit
- NPL and EPL estimators approach MLE

Monte Carlo B: discontinuous likelihood

Number of equilibria at true parameter: 9
Number of equilibria in the data: 1
Data generating equilibrium: unstable, near "cliffs"

	2step	NPL	EPL	MPEC-VP	MPEC-P	NRLS
True k1=3.5	3.49739	3.55144	3.64772	3.65943	3.67027	3.50212
Bias	-0.00261	0.05144	0.14772	0.15943	0.17027	0.00212
MCSD	0.13999	0.07133	0.12900	0.12693	0.11583	0.03255
ave log-like	-0.27494	-0.29474	-0.29528	-0.30330	-0.30257	-0.25086
log-likelihood	-1374.721	-1473.695	-1476.425	-1516.503	-1512.847	-1254.320
log-like short	-	-219.375	-222.104	-270.999	-267.523	-0.000
KL divergence	0.01512	0.04889	0.04495	0.04102	0.04078	0.00016
$\|\|P - P_{\mathbf{0}}\|\|$	0.62850	0.86124	0.83062	0.66562	0.65879	0.01610
$\|\|\Psi(P) - P\|\|$	0.763764	0.000000	0.000000	0.000000	0.000000	0.000002
$\|\|\Gamma(v) - v\|\|$	0.852850	0.000000	0.000000	0.000000	0.000000	0.000005
N runs of 100	100	100	100	28	27	100

- Equilibrium conditions are satisfied, but estimators converge to *wrong* equilibria as seen from KL divergence from DGP equilibria
- Biased estimates by EPL, NPL and MPEC
(constraints are satisfied, yet low likelihood and high KL divergence)

Likelihood correspondence

Lines are costructed using symmetric KL-divergence

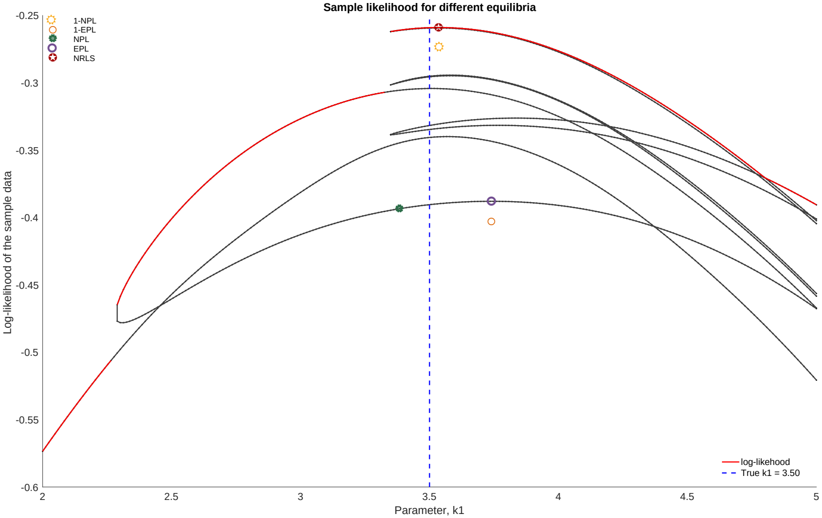

Equilibrium selection estimates

Distribution of KL-divergence to the DGP equilibrium, vertical lines represent other equilibria

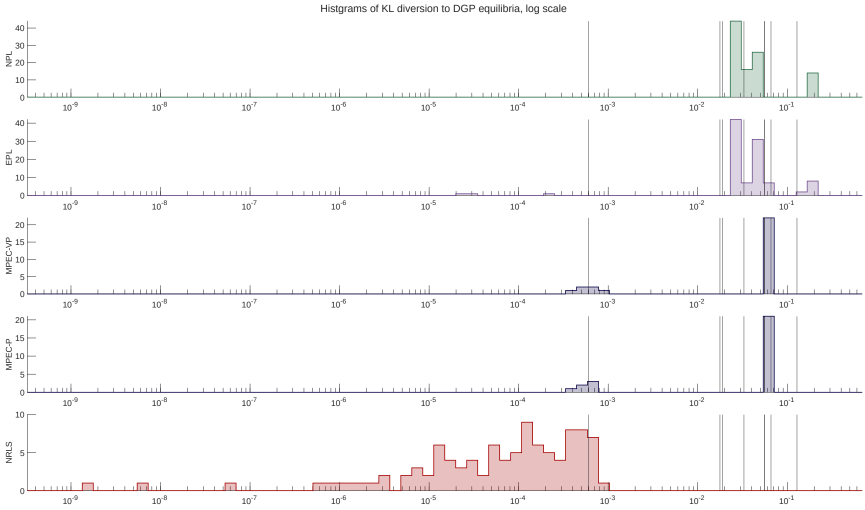

Monte Carlo B: massive multiplicity

Number of equilibria at true parameter: 2455

Number of equilibria in the data: 1

Time to enumerate all equilibria (RLS) once: 10m 39s

	1-NPL	NPL	EPL	NRLS
True k1=3.75	3.70959	3.71272	3.78905	3.74241
Bias	-0.04041	-0.03728	0.03905	-0.00759
MCSD	0.11089	0.06814	0.40716	0.03032
ave log-likelihood	-0.38681557	-0.37348793	-0.45256293	-0.35998461
log-likelihood	-1934.078	-1867.440	-2262.815	-1799.923
log-like shortfall	-	-66.529	-467.607	-0.000
KL divergence	Inf	14.07523	12231.59186	0.32429
$\|\|P - P_{\mathbf{0}}\|\|$	0.82204	0.65580	0.79241	0.07454
$\|\|\Psi(P) - P\|\|$	0.963574	0.000000	0.000000	0.000006
$\|\|\Gamma(v) - v\|\|$	7.020899	0.000000	0.000000	0.000008
N runs of 100	100	18	68	100
CPU time	0.159s	11.262s	4.013s	4.731s

- Severe convergence problems for NPL and EPL
- Poor eqb identification (low likelihood and high KL divergence)
- NRLS has comparable CPU time (much faster than full enumeration)

Monte Carlo C, multiple equilibria in the data

- Assume that the same equilibrium is played in each market over time
- Grouped fixed-effects, groups defined by the equilibria played

1. Joint grouped fixed-effects estimation
    - Estimate the partition of the markets into groups playing different equilibria together with θ
    - For each market compute maximum likelihood over all equilibria and "assign" it to the relevant group (estimation+classification)
    - Computationally very demanding: BnB market-by-market, non-parametric refinement has no bite
2. Two-step grouped fixed-effects estimation
    - Step 1: partition the markets based on some observable characteristics (K-means clustering) (Outside of Monte Carlo)
    - Step 2: estimate θ allowing different equilibria in different groups
    - Small additional computational cost for NRLS!

Bonhomme, Manresa (2015); Bonhomme, Lamadon, Manresa (2022)

Monte Carlo C: multiple equilibria in the data

Number of equilibria at true parameter: 81
Number of equilibria in the data: 5
Number of unique equilibria in the data: 3

	1-NPL	NRLS				
True k1=9.25	9.20991	9.25449				
Bias	-0.04009	0.00449				
MCSD	0.15021	0.04109				
ave log-likelihood	-0.798223	-0.707174				
log-likelihood	-19955.57	-17679.36				
log-like shortfall	-	0.000				
KL divergence	0.32943	0.00039	0.00039	0.00039	0.00040	0.00028
$\|\|P - P_{\mathbf{0}}\|\|$	0.32787	0.00287	0.00287	0.00287	0.00252	0.00240
$\|\|\Psi(P) - P\|\|$	0.460870	0.000000	0.000000	0.000000	0.000000	0.000000
$\|\|\text{Bellman}(V) - V\|\|$	5.438776	0.000000	0.000000	0.000000	0.000000	0.000000
# converged of 100	100	100				
CPU time, sec	0.023	20.695				

- All 5 equilibria were identified correctly as seen from KL divergence
- The first three equilibria are the same in DGP, and have the same KL and L1 divergence
- Similar results in runs with many more equilibria in the data

Equilibrium selection estimates

Distribution of KL-divergence to the DGP equilibrium, vertical lines represent other equilibria

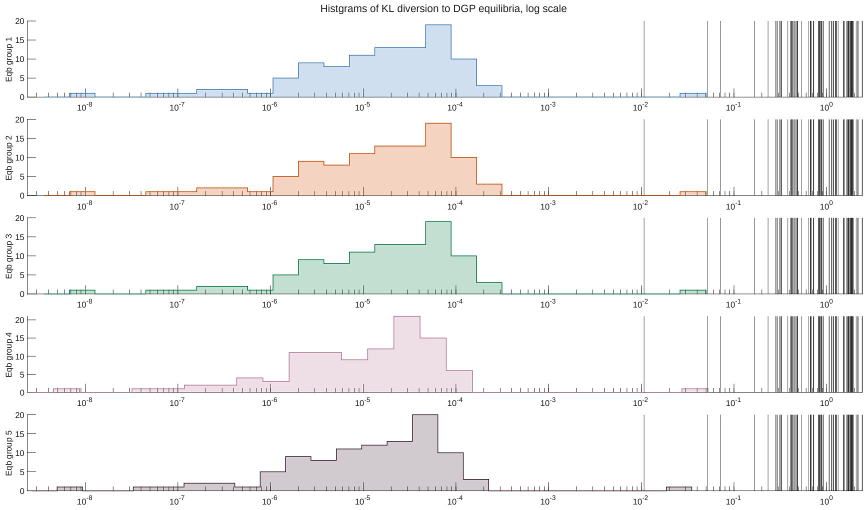

Monte Carlo C, run 2: many more equilibria in the data

Number of equilibria at true parameter: 19,683
Number of equilibria in the data: 25

	1-NPL	EPL	NRLS
True k1=9.5	9.12069	9.40900	9.49992
Bias	0.03918	-0.09100	-0.00008
MCSD	0.02864	0.05586	0.00414
ave log-likelihood	-0.511188	-0.515356	-0.504482
log-likelihood	-25559.42	-25767.79	-25224.10
log-like shortfall	-	-543.690	0.000
KL divergence	0.05753	0.05045	0.0 to 0.0001
$\|\|P - P_{\mathbf{0}}\|\|$	0.18517	0.26430	0.00059 to 0.00572
$\|\|\Psi(P) - P\|\|$	0.186981	0.000000	0.0 for all
$\|\|\mathrm{Bellman}(V) - V\|\|$	2.577006	0.000000	0.0 for all
# converged of 100	100	100	100
CPU time, sec	0.047	1.041	3m 9.4s

- All 25 equilibria were identified correctly
- Largest average KL divergence 10^{-4} whereas the closest to DGP equilibrium at true θ has KL=10^{-2}
- Largest error in choice probabilities across the state space 0.00572

Equilibrium selection estimates

Distribution of KL-divergence to the DGP equilibrium, vertical lines represent other equilibria

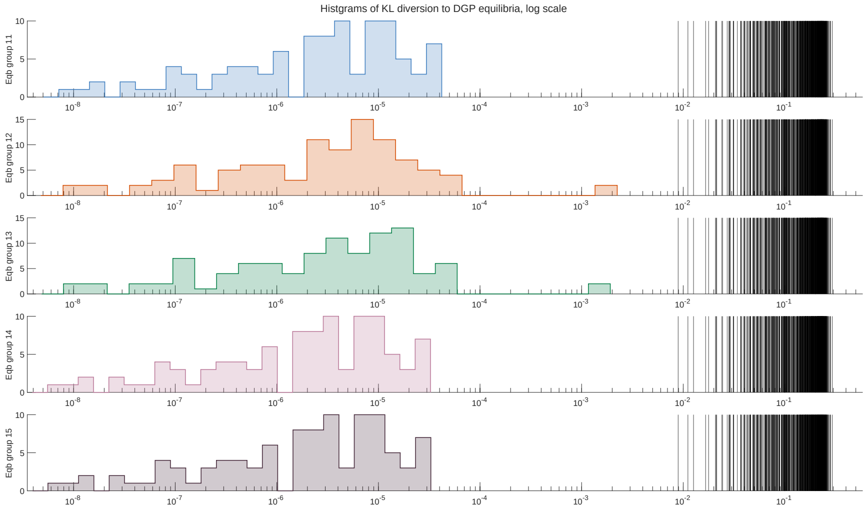

Shown are the first five identified equilibria, other twenty are similar

NRLS estimator for directional dynamic games

Complicated computational task involving maximization over the large finite set of all MPE equilibria $\rightarrow$ branch-and-bound algorithm with combined likelihood bounding rule

1. Each stage game $\rightarrow$ non-linear solver, specific to the model
2. Combining stage game solutions to full game MPEs $\rightarrow$ State Recursion algorithm
3. Solving for all MPE equilibria $\rightarrow$ Recursive Lexicographic Search
4. Structural estimation $\rightarrow$ Nested Recursive Lexicographic Search

- Implementation of statistically efficient estimator (MLE)
- Using BnB NRLS avoids full enumeration at no cost
- BnB augmented with non-parametric likelihood bounding function $\rightarrow$ less computation with larger sample size
- Computationally tractable
- Fully robust to multiplicity of equilibria