


Structural Estimation of Directional Dynamic Games With Multiple Equilibria

Fedor Iskhakov, Stony Brook University
Dennis Kristensen, University College London
John Rust, Georgetown University
Bertel Schjerning, University of Copenhagen

Econometrics seminar
Department of Economics
Ohio State University
October 5, 2026

Simultaneous move leapfrogging: 164,295,079 equilibria



Iskhakov, Rust and Schjerning (2016) *Review of Economic Studies*
“Recursive Lexicographical Search: Finding All Markov Perfect Equilibria
of Finite State Directional Dynamic Games”

Structural estimation of stochastic dynamic games

Structural model that includes:

1. Several decision makers (*players*)
 2. Maximize discounted expected lifetime utility
 3. Anticipate consequences of their current actions
 4. Anticipate actions by other players in current and future periods (*strategic interaction*)
 5. Operate in a stochastic environment (*state of the game*) whose evolution depend on the collective actions of the players
- ▶ Data on $M \rightarrow \infty$ independent markets over $T < \infty$ periods
 - ▶ Observed are states and actions of all players
 - ▶ Multiplicity of equilibria in both the model and potentially the data
 - ▶ Nested MLE estimator based on full solution of the model

Markov Perfect Equilibria

- ▶ Discrete-time infinite-horizon dynamic stochastic games with discrete states and actions
- ▶ MPE is a pair of **strategy profiles** and **value functions** such that

$$V = \Psi^V(V, P, \theta) \quad (\text{Bellman equations})$$

$$P = \Psi^P(V, P, \theta) \quad (\text{CCPs} = \text{mutual best responses})$$

- ▶ $\Psi = (\Psi^V, \Psi^P)$ gives the structure of the model
- ▶ With Hotz-Miller inversion can be reformulated only in P or in V
- ▶ Denote the set of all equilibrium choice probabilities in the model as

$$\mathcal{E}(\Psi, \theta) = \left\{ P \mid \exists V : \begin{array}{l} V = \Psi^V(V, P, \theta) \\ P = \Psi^P(V, P, \theta) \end{array} \right\}$$

- ▶ Approach: Solve for all MPE equilibria for any θ

Maximum Likelihood Estimation

- ▶ Data from M independent markets from T periods

$$Z = \{\bar{a}^{mt}, \bar{x}^{mt}\}_{m \in \mathcal{M}, t \in \mathcal{T}}$$

- ▶ Assume that only one equilibrium is played in the data (we relax this assumption later $\rightarrow$ grouped fixed effects)
- ▶ For a given θ denote the choice probabilities for player i at time t and market m as $P_i(a_i^{mt} | x^{mt}; \theta)$

$$P(\theta) \in \mathcal{E}(\Psi, \theta) : P(\theta) = \left\{ P_i(a_i^{mt} | x^{mt}; \theta) \right\}_{i,m,t}$$

- ▶ MLE estimator $\hat{\theta}^{ML}$ is given by

$$\hat{\theta}^{ML} = \arg \max_{\theta} \left[\max_{P(\theta) \in \mathcal{E}(\Psi, \theta)} \frac{1}{M} \sum_{i=1}^N \sum_{m=1}^M \sum_{t=1}^T \log P_i(\bar{a}_i^{mt} | \bar{x}^{mt}; \theta) \right]$$

MLE via Constrained Optimization Approach

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

$$\mathcal{L}(Z, P, \theta) = \frac{1}{M} \sum_{i=1}^N \sum_{m=1}^M \sum_{t=1}^T \log P_i(\bar{a}_i^{mt} | \bar{x}^{mt}; \theta)$$

- ▶ The constrained optimization formulation of the ML estimation problem is

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

- ▶ **Math programming with equilibrium constraints (MPEC)**
- ▶ Does not rely *as much* on the structure of the problem
- ▶ Much bigger computational problem
- ▶ Implements the same MLE estimator (*when it works*)



Su (2014); Egesdal, Lai and Su (2015)

Estimation methods for dynamic stochastic games

▶ Two step (CCP) estimators

- ▶ Fast, do not impose equilibrium constraints, finite sample bias

1. Estimate CCP $\rightarrow \hat{P}$
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 (2025), ReStud

Overview of NRLS

Full solution nested fixed point MLE estimator

with computational enhancements to ensure tractability

- ▶ Reliable and *computationally feasible*^(?) MLE estimator for **directional dynamic games (DDG)**
- ▶ Rely of full solution algorithm
- ▶ Discrete programming methods (BnB) to economize on compute
- ▶ Use non-parametric likelihood for effective bounding
- ▶ **Fully reliable in presence of multiplicity of MPE**
- ▶ **Relax single-equilibrium-in-data assumption**

Asymptotic properties of the estimator — in a companion paper

- ▶ Topology of the equilibrium/likelihood correspondence
- ▶ Consistency and asymptotic normality of the estimator
- ▶ Convergence rate for the equilibrium identification

ROAD MAP

1. Solving directional dynamic games (DDGs):
 - ▶ Bertrand pricing and investment game
 - ▶ Solving once: State recursion algorithm
 - ▶ Solving for all MPE: Recursive lexicographical search (RLS) algorithm
2. Structural estimation of DDGs: nested MLE RLS
 - ▶ Is it even feasible computationally?
3. Monte Carlo to compare to existing estimators in dynamic games (two-step CCP, NPL, EPL, MPEC)
 - ▶ One equilibrium in the model and data
 - ▶ Multiplicity of equilibria at true parameter
 - ▶ Multiple equilibria in the data
4. Application to the McDonald's vs. Burger King entry game
 - ▶ Data from Toivanen and Waterson (2005)



Iskhakov, Rust and Schjerning, 2016, ReStud (full solution)
Iskhakov, Rust and Schjerning, 2018, IER (leapfrogging)

Amcor-Visy collusion case



- ▶ Australian market for **cardboard** is essentially a **duopoly**
- ▶ Between 2000 and 2005 *Visy* and *Amcor* **colluded** to divide the market of cardboard and to fix prices
- ▶ 2007: **Visy admits** to have been manipulating the market, issued with \$36 million fine
- ▶ July 2009: Cadbury vs. Amcor, **damages estimated at \$235.8 million**, settles out of court
- ▶ March 2011: **Class action suit** against both Amcor and Visy settles out of court for \$95 million

Cardboard industry in Australia

- ▶ Cardboard is a highly **standardized product**
- ▶ **Bertrand price competition** with strong incentives for price cutting
- ▶ Amcor and Visy do **minimal amounts of R&D** themselves,
- ▶ Instead **purchase new technology from other companies** to **reduce cost of production**
 - ▶ Amcor plans to build state-of-the-art paper mill in Botany Bay **before** the collusion took place
 - ▶ “B9” plant finally opened on **February 1, 2013**

Leapfrogging equilibrium

- ▶ Firms invest in alternating fashion and take turns in cost leadership
- ▶ Market price makes permanent downward shifts

Dynamic Bertrand price competition

Directional stochastic dynamic game

- ▶ Two Bertrand competitors, $n = 2$, no entry or exit
- ▶ Discrete time, infinite horizon ($t = 1, 2, \dots, \infty$)
- ▶ Firms maximize expected discounted profits
- ▶ Each firm has two choices in each period:
 1. Price for the product — simultaneous
 2. Whether or not to buy the state of the art technology
 - ▶ Simultaneous moves
 - ▶ Alternating moves

Static Bertrand price competition in each period

- ▶ Continuum of consumers make static purchase decision
- ▶ No switching costs: buy from the lower price supplier
- ▶ Per period profits (c_i is the marginal cost)

$$r_i(c_1, c_2) = \begin{cases} 0 & \text{if } c_i \geq c_j \\ c_j - c_i & \text{if } c_i < c_j \end{cases}$$

Cost-reducing investments

State-of-the-art production cost c process

- ▶ Initial value c_0 , lowest value 0: $0 \leq c \leq c_0$
- ▶ Discretized with n points
- ▶ Follows exogenous Markov process and only improves
- ▶ Markov transition probability $\pi(c_{t+1}|c_t)$
 $\pi(c_{t+1}|c_t) = 0$ if $c_{t+1} > c_t$

State space of the problem

- ▶ State of the game: cost structure (c_1, c_2, c)
- ▶ State space is $S = (c_1, c_2, c) \subset R^3$: $c_1 \geq c$, $c_2 \geq c$
- ▶ Actions are observable
- ▶ Private information EV(1) i.i.d. shocks $\eta\epsilon_{i,I}$ and $\eta\epsilon_{i,N}$

Bellman equations, firm 1, simultaneous moves

$$V_1(c_1, c_2, c, \epsilon_1) = \max [v_1(I, c_1, c_2, c) + \eta\epsilon_1(I), v_1(N, c_1, c_2, c) + \eta\epsilon_1(N)]$$

$$\begin{aligned}v_1(N, c_1, c_2, c) &= r_1(c_1, c_2) + \beta EV_1(c_1, c_2, c, N) \\v_1(I, c_1, c_2, c) &= r_1(c_1, c_2) - K(c) + \beta EV_1(c_1, c_2, c, I)\end{aligned}$$

With extreme value shocks, the investment probability (CCP) is

$$P_1(I|c_1, c_2, c) = \frac{\exp\{v_1(I, c_1, c_2, c)/\eta\}}{\exp\{v_1(I, c_1, c_2, c)/\eta\} + \exp\{v_1(N, c_1, c_2, c)/\eta\}}$$

- There is a separate Bellman equation for player 2, with “outputs” V_2 and P_2 , where $P_2(I|c_1, c_2, c)$ is firm 2’s probability of investing in state (c_1, c_2, c) .

Bellman equations, firm 1, simultaneous moves

The expected values are given by

$$EV_1(c_1, c_2, c, N) = \int_0^c \left[P_2(I|c_1, c_2, c) H_1(c_1, c, c') + [1 - P_2(I|c_1, c_2, c)] H_1(c_1, c_2, c') \right] \pi(dc'|c)$$


$$EV_1(c_1, c_2, c, I) = \int_0^c \left[P_2(I|c_1, c_2, c) H_1(c, c, c') + [1 - P_2(I|c_1, c_2, c)] H_1(c, c_2, c') \right] \pi(dc'|c)$$

$$H_1(c_1, c_2, c) = \eta \log \left[\exp(v_1^N(c_1, c_2, c)/\eta) + \exp(v_1^I(c_1, c_2, c)/\eta) \right].$$

is the “smoothed max” or logsum function





Discretized state space = a “quarter pyramid”

$$S = \{(c_1, c_2, c) | c_1 \geq c, c_2 \geq c, c \in [0, 3]\}, n = 4$$




Transitions due to technological progress

As c decreases, the game falls through the layers of the pyramid




Game dynamics: example

The game starts at the apex, as some point technology improves




Game dynamics: example

Both firms buy new technology $c = 2 \rightsquigarrow (c_1, c_2, c) = (2, 2, 2)$




Game dynamics: example

State-of-the-art technology becomes $c = 1 \rightsquigarrow (c_1, c_2, c) = (2, 2, 1)$




Game dynamics: example

Firm 1 invests and becomes cost leader $\rightsquigarrow (c_1, c_2, c) = (1, 2, 1)$




Game dynamics: example

State-of-the-art technology becomes $c = 0 \rightsquigarrow (c_1, c_2, c) = (1, 2, 0)$




Game dynamics: example

A particular sequence of investment decisions along technological progress pass




Transitions due to technological progress

As c decreases, the game falls through the layers of the pyramid




Strategy-specific partial order on S

Strategy σ_1 of firm 1: invest at all interior points




Strategy-specific partial order on S

Strategy σ_2 of firm 2: invest at all edge points




Strategy-specific partial order on S

Strategy $\sigma = (\sigma_1, \sigma_2)$ of both firms



Strategy independent partial order on S


Coarsest common refinement of partial orders induced by all strategies




Definition of the Dynamic Directional Games

Finite state Markovian stochastic game is a DDG if it holds:

1. Every feasible Markovian strategy σ satisfies the no loop condition.




2. Every pair of feasible Markovian strategies σ and σ' induce consistent partial orders on the state space.



Total order on the set of stages

After running a topological sort algorithm on the DAG


DAG toposort



State recursion algorithm



Stage recursion

Backward induction on stages of DDG instead of time



Examples of Directional Dynamic Games

Many games have state dynamic evolutions described by a DAGs



Judd, Schmedders, Yeltekin (2012), *IER*


“Optimal rules for patent researchers”



Dube, Hitsch, Chintagunta (2010), *Marketing Science*

“Tipping and concentration in markets with indirect network effects”

Tennis is a Directional Dynamic Game




Anderson, Rosen, Rust, Wong (2025) *JPE*
“Disequilibrium Play in Tennis”

Multiplicity of stage equilibria

Number of equilibria in the higher stages depends on the selected equilibria

- ▶ State recursion proceeds **conditional on** equilibrium selection rule
- ▶ Multiplicity of stage equilibria $\Leftrightarrow$ multiplicity
- ▶ Can systematically combine different stage equilibria




Best response maps


Typically one or three stage equilibria, up to 5

When $\eta \rightarrow 0$

Correspondences for $\eta = 0$



► Smooth functions with $\eta > 0$



Recursive Lexicographic Search Algorithm

Building blocks of RLS algorithm:

1. State recursion algorithm solves the game **conditional on** equilibrium selection rule (ESR)
2. RLS algorithm efficiently cycles through **all feasible** ESRs

Challenge:


- ▶ Choice of a particular MPE for any stage game at any stage
- ▶ may alter the **set** and even the **number** of stage equilibria at earlier stages

Solution: **RLS = depth-first tree traversal** (illustration coming)


- ▶ Root of the tree is one of the absorbing states
- ▶ Levels of the tree correspond to the state points
- ▶ Branching happens when stages have multiple equilibria
- ▶ MPE of the game is given by a path from root to a leaf

Indexing to preserve the directional dependence


End game points have the highest index, root has the lowest index




RLS Tree Traversal, step 1




RLS Tree Traversal, step 2




RLS Tree Traversal, step 3




RLS Tree Traversal, step 4




RLS Tree Traversal, step 5




RLS Tree Traversal, step 6




RLS Tree Traversal, step 7



RLS Tree Traversal, step 8



RLS Tree Traversal, step 9



Recursive Lexicographic Search (RLS) algorithm

Theorem (RLS theorem)

Assume there exists an algorithm that can find all MPE of every stage game of the DDG, and that the number of these equilibria is finite in every stage game.

Then the RLS algorithm finds all MPE of the DDG in a finite number of steps, which equals the total number of MPE.



Iskhakov, Rust and Schjerning, 2016, ReStud

“Recursive lexicographical search: Finding all markov perfect equilibria of finite state directional dynamic games.”

ROAD MAP

1. Solving directional dynamic games (DDGs):
 - ▶ Bertrand pricing and investment game
 - ▶ Solving once: State recursion algorithm
 - ▶ Solving for all MPE: Recursive lexicographical search (RLS) algorithm
2. Structural estimation of DDGs: nested MLE RLS
 - ▶ Is it even feasible computationally?
3. Monte Carlo to compare to existing estimators in dynamic games (two-step CCP, NPL, EPL, MPEC)
 - ▶ One equilibrium in the model and data
 - ▶ Multiplicity of equilibria at true parameter
 - ▶ Multiple equilibria in the data
4. Application to the McDonald's vs. Burger King entry game
 - ▶ Data from Toivanen and Waterson (2005)

Nested Recursive Lexicographical Search

- ▶ Data from M independent markets from T periods

$$Z = \{\bar{a}^{mt}, \bar{x}^{mt}\}_{m \in \{1, \dots, M\}, t \in \{1, \dots, T\}}$$

- ▶ Let the set of all MPE choice probabilities be

$$\mathcal{E}(\Psi, \theta) = \{P_{\theta}^1, \dots, P_{\theta}^{K(\theta)}\}$$

1. Outer loop

Maximization of the likelihood function w.r.t. to **structural parameters** 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

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

Max of a function on a discrete set organized into RLS tree

Likelihood over the state space

- ▶ Given **equilibrium** k choice probabilities
 $P_\theta^k(a|x) = (P_1^k(a|x, \theta), \dots, P_N^k(a|x, \theta))$, likelihood is

$$\mathcal{L}(Z, \theta, P_\theta^k) = \frac{1}{M} \sum_{m=1}^M \sum_{t=1}^T \sum_{i=1}^N \log P_i^k(\bar{a}_i^{mt} | \bar{x}^{mt}; \theta)$$

- ▶ Let ι **index points in the state space**, $\iota \in \{1, \dots, S\}$
 $\iota = 1$ initial point, $\iota = S$ the terminal (absorbing) state
- ▶ Denote n_ι the number of observations in state x_ι and $n_\iota^{a_i}$ the number of observations of player i taking action a_i at x_ι


$$n_\iota = \sum_{m=1}^M \sum_{t=1}^T \mathbb{1}\{\bar{x}^{mt} = x_\iota\} \quad n_\iota^{a_i} = \sum_{m=1}^M \sum_{t=1}^T \mathbb{1}\{\bar{a}_i^{mt} = a_i, \bar{x}^{mt} = x_\iota\}$$

- ▶ Then equilibrium-specific likelihood can be computed as

$$\mathcal{L}(Z, \theta, P_\theta^k) = \frac{1}{M} \sum_{\iota=1}^S \sum_{i=1}^N \sum_a n_\iota^{a_i} \log P_i^k(a|x_\iota; \theta)$$

Data distribution over the state space

1000 markets, 5 time periods, init at apex of the pyramid



Branch and bound (BnB) method

Old method for solving integer programming problems

1. Form a **tree** of subdivisions of the set of admissible plans
 $\implies$ **RLS tree**
2. Specify a **bounding function** representing the best attainable objective on a given subset (branch)
 $\implies$ **Partial likelihood function** computed on the subset of states

$$\mathcal{L}^{\text{part}}(\mathbf{Z}^{\mathcal{S}}, \theta, V_{\theta}^k) = \sum_{\iota \in \mathcal{S}} \sum_{i=1}^N \sum_a n_{\iota}^{a_i} \log P_i^k(a|x_{\iota}; \theta)$$

where $\mathbf{Z}^{\mathcal{S}} = \{(a, x) : x \in \mathcal{S} \subset \{1, \dots, S\}\}$ denotes data 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 $\mathbf{Z}^{\mathcal{S}} = \mathbf{Z}$
(at the leaf of RLS tree)
3. Dismiss the subsets of the plans where the bound is below the current best attained value of the objective




Land and Doig, 1960 *Econometrica*


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 33



BnB on RLS tree, step 34



Non-parametric likelihood bounding

- ▶ **Idea**: augment the partial likelihood bounding function with non-parametric likelihood computed on the “rest of the data”
- ▶ Parametric choice probabilities $P_i^k(a|x_\ell; \theta) \rightarrow$ frequencies $n_\ell^{a_i}/n_\ell$

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

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

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


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

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

- ▶ **Augmented partial likelihood is much more powerful bound for BnB**


Non-parametric likelihood bounding

$\iota = S = 14$ (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 $\rightarrow$ parametric likelihood as $N \rightarrow \infty$ at true $\theta \Rightarrow$ even less computation



Full enumeration RLS in larger sample

Comparing with the previous slide most of the computation is avoided!



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
- ▶ With more data as $M \rightarrow \infty$
- ▶ Non-parametric log-likelihood converges to the true partial likelihood (assuming the model is well-specified and in the neighborhood of the true parameter)
- ▶ The width of the band between the blue lines in the plots decreases
 - $\rightarrow$ Even sharper Bounding Rules
 - $\rightarrow$ Even less computation
 - $\rightarrow$ Still delivering exact maximum

MLE for any sample size, but much easier to compute with more data!

ROAD MAP

1. Solving directional dynamic games (DDGs):
 - ▶ Simple example: Bertrand pricing and investment game
 - ▶ State recursion algorithm
 - ▶ Recursive lexicographical search (RLS) algorithm
2. Structural estimation of DDGs using Nested RLS
 - ▶ Branch-and-bound on RLS tree
 - ▶ Non-parametric likelihood bounding
3. Monte Carlo: (Compare NRLS, two-step CCP, NPL, EPL, MPEC)
 - ▶ One equilibrium in the model and data
 - ▶ Multiplicity of equilibria at true parameter
 - ▶ (Multiple equilibria in the data)
4. Application to the McDonald's vs. Burger King entry game
 - ▶ Data from Toivanen and Waterson (2005)

Monte Carlo simulations

A

Single equilibrium in the model
One equilibrium in the data

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

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

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 - P_0 $	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, run 1: little multiplicity

Number of equilibria at true parameter: 3

Number of equilibria in the data: 1

Data generating equilibrium: [stable](#)

	2step	NPL	EPL	MPEC-VP	MPEC-P	NRLS
True k1=7.5	7.55163	7.49844	7.49918	7.65318	7.35124	7.49919
Bias	0.05163	-0.00156	-0.00082	0.15318	-0.14876	-0.00081
MCS D	0.17875	0.06062	0.03413	0.99742	0.47136	0.03413
ave log-likelihood	-0.84779	-0.84425	-0.84421	-0.88682	-0.87541	-0.84421
log-likelihood	-21194.86	-21106.33	-21105.13	-22170.40	-21885.37	-21105.13
log-likelihood short	-	-1.206	-0.000	-1062.740	-776.809	-0.000
KL divergence	0.02557	0.00040	0.00013	0.23536	0.16051	0.00013
$\ P - P_0\ $	0.11085	0.00490	0.00280	0.17466	0.20957	0.00280
$\ \Psi(P) - P\ $	0.170940	0.000000	0.000000	0.000000	0.000000	0.000000
$\ \Gamma(v) - v\ $	1.189853	0.000000	0.000000	0.000000	0.000000	0.000001
N runs of 100	100	100	100	98	97	100

- ▶ MPEC convergence deteriorates
- ▶ Equilibrium conditions are satisfied, but estimators start to converge to *wrong* equilibria
(as seen from KL divergence from the data generating equilibrium)

Monte Carlo B, run 2: little multiplicity, unstable

Number of equilibria at true parameter: 3

Number of equilibria in the data: 1

Data generating equilibrium: **unstable**

	2step	NPL	EPL	MPEC-VP	MPEC-P	NRLS
True $k_1=7.5$	7.54238	7.39276	7.48044	7.73133	7.63100	7.50176
Bias	0.04238	-0.10724	-0.01956	0.23133	0.13100	0.00176
MCSD	0.17145	0.05608	0.15801	0.72988	0.89874	0.03820
ave log-like	-0.86834	-0.89374	-0.86550	-0.88512	-0.90196	-0.86504
log-likelihood	-21708.60	-22343.58	-21637.54	-22127.91	-22549.06	-21626.12
log-like short	-	-765.242	-11.413	-502.121	-920.643	-0.000
KL divergence	0.02271	0.15996	0.00257	0.11452	0.20182	0.00012
$\ P - P_0\ $	0.09757	0.20709	0.00619	0.03860	0.02504	0.00307
$\ \Psi(P) - P\ $	0.160102	0.000000	0.000000	0.000000	0.000000	0.000000
$\ \Gamma(v) - v\ $	1.126738	0.000000	0.000000	0.000000	0.000000	0.000001
N runs of 100	100	18	100	99	98	100


- ▶ NPL estimator fails to converge
- ▶ Similar convergence issues for MPEC
- ▶ EPL estimator performs well



Aguirregabiria, Marcoux (2021)

Likelihood correspondence

Lines are constructed using symmetric KL-divergence



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 $k_1=3.5$	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
MCS D	0.13999	0.07133	0.12900	0.12693	0.11583	0.03255
ave log-likelihood	-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-likelihood 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_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 constructed 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 $k_1=3.75$	3.70959	3.71272	3.78905	3.74241
Bias	-0.04041	-0.03728	0.03905	-0.00759
MCS D	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_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: NRLS as the sample grows

Vary number of markets M , fix $T = 5$ (asymptotics of the theory)

Smooth likelihood (MC B, run 1)


- ▶ True $k_1 = 7.5$
- ▶ 3 MPEs at true parameter
- ▶ Data from the **stable** equilibrium
- ▶ $M \in \{50, 75, 100, 150, 200, 250, 350, 500, 750, 1000, 1250, 1500, 1750, 2000, 2250, 2500, 3500, 5000\}$
- ▶ 300 replications at every M , **1,000** at $M \in \{50, 250, 1000, 2500, 5000\}$
- ▶ Samples nested across M within a replication
- ▶ NRLS initialized at 1-NPL; estimate audited by exact RLS on a grid
- ▶ “Other MPE”: selected eqb is not the continuation of the DGP branch to $\hat{k}_1$

Discontinuous likelihood (MC B)

- ▶ True $k_1 = 3.5$
- ▶ 9 MPEs at true parameter
- ▶ Data from an **unstable** eqb near the likelihood cliff

Statistical convergence of NRLS


M	Smooth, $k_1 = 7.5$			Discontinuous, $k_1 = 3.5$		
	RMSE	$\sqrt{M}$ RMSE	Other MPE	RMSE	$\sqrt{M}$ RMSE	Other MPE
50	0.3071	2.17	2.2%	0.1405	0.99	37.2%
250	0.1397	2.21	0.0%	0.0654	1.03	18.7%
1000	0.0724	2.29	0.0%	0.0350	1.11	13.5%
2500	0.0468	2.34	0.0%	0.0228	1.14	5.9%
5000	0.0317	2.24	0.0%	0.0155	1.09	1.0%



- ▶ $\sqrt{M}$ RMSE flat $\Rightarrow$ root- M rate in both designs; bias small relative to RMSE
- ▶ Near the cliff wrong-branch selection is frequent in small samples, but vanishes with M

Asymptotic normality of NRLS


$\sqrt{M}(\hat{k}_1 - k_1^0)$ over 1,000 replications; dashed: normal with variance at $M = 5000$



- ▶ Smooth likelihood: close to normal already at $M = 50$
- ▶ Cliff: mass point from **other-MPE** selections at $M = 50$ (37.2%), gone by $M = 5000$ (1.0%)
- ▶ Finite-sample branch switching $\neq$ nonstandard asymptotics: θ_0 is away from the boundary of the DGP eqb existence set

Computational cost falls with sample size

Median NRLS time (10th–90th percentile), MATLAB R2025a, Apple M4 Max



	Smooth		Discontinuous	
	$M = 50$	$M = 5000$	$M = 50$	$M = 5000$
Median NRLS time, sec	2.43	0.16	0.64	0.25
Likelihood evaluations (mean)	22.3	15.1	20.0	10.9
Extended equilibria at $\hat{k}_1$ (mean)	1.26	1.00	2.29	1.15
Stage games solved (mean)	4,805	484	3,785	1,012

- More data $\Rightarrow$ sharper likelihood bound $\Rightarrow$ BnB prunes more of the RLS tree

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 $k_1=9.25$	9.20991	9.25449				
Bias	-0.04009	0.00449				
MCS D	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_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 $k_1=9.5$	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_0\ $	0.18517	0.26430	0.00059 to 0.00572
$\ \Psi(P) - P\ $	0.186981	0.000000	0.0 for all
$\ \text{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 $\text{KL}=10^{-2}$
- ▶ Largest error in choice probabilities across the state space 0.00572

ROAD MAP

1. Solving directional dynamic games (DDGs):
 - ▶ Simple example: Bertrand pricing and investment game
 - ▶ State recursion algorithm
 - ▶ Recursive lexicographical search (RLS) algorithm
2. Structural estimation of DDGs using Nested RLS
 - ▶ Branch-and-bound on RLS tree
 - ▶ Non-parametric likelihood bounding
3. Monte Carlo: (Compare NRLS, two-step CCP, NPL, EPL, MPEC)
 - ▶ One equilibrium in the model and data
 - ▶ Multiplicity of equilibria at true parameter
 - ▶ (Multiple equilibria in the data)
4. Application to the McDonald's vs. Burger King entry game
 - ▶ Data from Toivanen and Waterson (2005)

Application: McDonald's vs. Burger King in the UK

Data

- ▶ Store stocks and openings of McDonald's (M) and Burger King (B) in 422 local authority districts (excl. Greater London), 1991–95
- ▶ 2,110 district-years, 4,220 binary entry choices
- ▶ Stores: M 281 $\rightarrow$ 421, B 79 $\rightarrow$ 181; districts with both: 55 $\rightarrow$ 121
- ▶ No closures: all 1,688 observed transitions weakly increase stocks

Why this application


- ▶ Irreversible entry $\Rightarrow$ game is **directional** by construction
- ▶ Established benchmark: NPL application of Aguirregabiria (2009), re-implemented by Hansen and Kofod Mogensen (2014), 15 parameters
- ▶ Does an explicit **likelihood maximum over all MPE** change the answer?



Toivanen and Waterson (2005) *RAND Journal of Economics*
Hansen and Kofod Mogensen (2014) University of Copenhagen
Aguirregabiria and Magesan (2020) *Review of Economic Studies*


Data: expansion of the two chains

Start-of-year chain presence by district



- **Expansion only:** districts move from *neither* to one chain to *both chains* (green spreads); none ever revert
- Districts with both chains: 55 in 1991 → 121 in 1995

Data: directionality in observed transitions



- All transitions on or above the diagonal: 1,485 of 1,688 are self-loops
- Only 38 of 225 states ever observed; max stocks B 7, M 13

Model: independent district entry games

- ▶ Players $i \in \{B, M\}$, action $a_i \in \{0, 1\}$: open at least one store
- ▶ State $x = (x_B, x_M) \in \{0, \dots, 14\}^2$, permanent district characteristics $z_m = (S_m, w_m)$: market size; density, GDP, rent, tax
- ▶ $N_i = x_i + a_i$, $x'_i = \min\{N_i, 14\}$, $D_i = N_i - N_{-i}$, flow payoff

$$\pi_i = \mathbb{1}\{N_i > 0\} \left[S_m (\theta_{i0}^{VP} + \theta_{i1}^{VP} D_i + \theta_{i2}^{VP} D_i^2) - (\theta_{i0}^{FC} + \theta_{i1}^{FC} N_i + \theta_{i2}^{FC} N_i^2) + N_i w'_m \gamma \right]$$

- ▶ Normal shocks, $\beta = 0.95$: $P_i(x, z_m) = \Phi(\Delta v_i(x, z_m; P, \theta))$, $P = \Psi(P; \theta)$

Directionality

- ▶ Stocks never decrease; only self-loop is “both wait” $\Rightarrow$ 225 states solved as a sequence of 2-player stage games
- ▶ z_m fixed: 422 disconnected blocks, $\mathcal{E}(\Psi, \theta) = \prod_m \mathcal{E}_m(\Psi, \theta)$

$$\mathcal{L}^{NRLS}(\theta) = \sum_{m=1}^{422} \max_{P_m \in \mathcal{E}_m(\Psi, \theta)} \mathcal{L}_m(\theta, P_m)$$

Estimates: 15-parameter model (preliminary)

		Reported NPL	Data 1-NPL	Replicated NPL	NRLS	NRLS s.e.
θ_{B0}^{VP}	B profit intercept	1.005	1.000508	1.004757	0.996945	0.197813
θ_{B1}^{VP}	B linear gap	-0.042	0.018537	-0.042306	-0.035191	0.028620
θ_2^{VP}	Shared quadratic gap	-0.008	0.000606	-0.007663	-0.007948	0.003364
θ_{B0}^{FC}	B cost intercept	0.070	0.088743	0.069751	0.087737	0.029140
θ_{B1}^{FC}	B linear cost	0.154	0.146583	0.153547	0.134218	0.023407
θ_{B2}^{FC}	B quadratic cost	-0.006	-0.005205	-0.005563	-0.003201	0.002031
θ_{M0}^{VP}	M profit intercept	1.425	1.264190	1.424467	1.303954	0.321717
θ_{M1}^{VP}	M linear gap	0.138	-0.013925	0.137983	0.181431	0.053768
θ_{M0}^{FC}	M cost intercept	0.088	0.096321	0.088362	0.079686	0.027776
θ_{M1}^{FC}	M linear cost	0.132	0.133309	0.131520	0.137539	0.021695
θ_{M2}^{FC}	M quadratic cost	-0.002	-0.005305	-0.001886	-0.001628	0.001277
γ_{dens}	Density	4.011	5.312279	4.010697	4.491863	1.332423
γ_{GDP}	GDP	0.118	0.177910	0.117513	0.135723	0.107355
γ_{rent}	Rent	-0.100	-0.078644	-0.100106	-0.083979	0.226232
γ_{tax}	Council tax	0.038	0.040553	0.037599	0.040636	0.026765
Log-likelihood		-898.680	—	-898.680	-897.864	

Data 1-NPL: pooled-probit first-stage CCPs (pseudo-LL = -893.87, not an equilibrium likelihood). NRLS s.e.: market-clustered outer product of scores on the locally followed branch, provisional. 4,220 choices, 422 districts.

- ▶ **NRLS raises the likelihood** above NPL ($-898.68 \rightarrow -897.86$); estimates broadly similar relative to s.e.; single MPE near the estimates

NRLS estimator for directional dynamic games

Complicated computational task involving maximization over the large finite set of all MPE equilibria → **branch-and-bound algorithm with refined bounding rule**

NRLS nested structure:


1. Each stage game → non-linear solver, **specific to the model**
2. Combining stage game solutions to full game MPEs → **State Recursion algorithm**
3. Solving for all MPE equilibria → **Recursive Lexicographic Search**
4. Structural estimation → **Nested loop MLE with BnB on RLS tree**

Performance of NRLS

- ▶ Implementation of statistically efficient estimator (MLE)
- ▶ BnB with NRLS avoids full enumeration at no additional cost
- ▶ BnB augmented with non-parameteric likelihood bound gives sharper Bounding Rules → less computation larger samples
- ▶ **Computationally tractable**
- ▶ **Fully robust to multiplicity of equilibria**


Slides are available for download at

<https://fedor.iskh.me>




DAG recursion to partition S into stages

Identify terminal states




DAG recursion to partition S into stages

Remove terminal states




DAG recursion to partition S into stages

Identify terminal states




DAG recursion to partition S into stages

Remove terminal states





DAG recursion to partition S into stages

Identify terminal states




DAG recursion to partition S into stages

Remove terminal states

[Back](#)


State recursion algorithm

Backward induction on stages of DDG





State recursion algorithm

Backward induction on stages of DDG




State recursion algorithm

Backward induction on stages of DDG




State recursion algorithm

Backward induction on stages of DDG




Backward induction on stages of DDG

Backward induction on stages of DDG




State recursion algorithm

Backward induction on stages of DDG




State recursion algorithm

Backward induction on stages of DDG




State recursion algorithm

Backward induction on stages of DDG




State recursion algorithm

Backward induction on stages of DDG



Best response functions as $\eta \rightarrow 0$



Asymptotic theory: setup and primitive properties

Setup as before:

- ▶ Players $i = 1..N$, finite states $s \in \mathcal{S}$ and actions $a_i \in \mathcal{A}_i(s)$, $\beta_i < 1$, stacked flow payoffs $U \in \mathbb{R}^{d_U}$, private shocks $\varepsilon_i(a_i)$
- ▶ CCPs $P \in \text{int } \mathcal{P}$ (dimension d_P), best response $\Phi(P; U) = \Lambda(\Delta v(P, U))$
- ▶ Equilibrium map $F(P; U) = P - \Phi(P; U)$, equilibrium CCPs $\mathcal{E}(U) = \{P \in \text{int } \mathcal{P} : F(P; U) = 0\}$

Definition $P \in \mathcal{E}(U)$ is **regular** if $\det D_P F(P; U) \neq 0$; $U \in \mathcal{U}_{\text{reg}}$ if all its equilibria are regular

Assumption 1 Finite $\mathcal{S}, \mathcal{A}$; shocks with full support and positive C^2 density, $E \max_a |\varepsilon_i(a)| < \infty$

Lemma 1 $F(P; U)$ is C^2 (calculus tools apply)

Lemma 2 Self-containment: on compact payoff sets equilibria stay in $\mathcal{P}_\epsilon = \{P \geq \epsilon\}$ (log-likelihood bounded)

Lemma 3 $\text{rank } D_U \Delta v(P, U) = d_P$: payoffs move choice values in every direction

Lemma 4 $\text{rank } D_U F = \text{rank } D_{(P,U)} F = d_P$ (total transversality)

Equilibrium correspondence $U \mapsto \mathcal{E}(U)$

Unrestricted payoffs U , regular equilibria:

Theorem Generic regularity: $\mathcal{U}_{\text{reg}}$ has full measure and is relatively open on compacts; each $U \in \mathcal{U}_{\text{reg}}$ has **finitely many** equilibria

Corollary Regular equilibrium $\Rightarrow$ unique local C^2 branch $P(U)$,
 $D_U P = -[D_P F]^{-1} D_U F$

Theorem Regular regions: on bounded connected open $\mathcal{U}_0 \subset \mathcal{U}_{\text{reg}}$ the equilibrium graph is a finite-sheeted covering: **constant number of globally labeled, uniformly separated C^2 branches** $P_e(U)$

Singular equilibria:

Proposition Everywhere: $\mathcal{E}(\cdot)$ nonempty, compact-valued, uhc; limits of equilibria are equilibria; isolated $\Rightarrow$ finite

Corollary Lyapunov–Schmidt nondegeneracy $\Rightarrow$ singular equilibrium is isolated

Corollary Folds: if $\text{rank } D_P F = d_P - 1$ with right/left null vectors q, r and $r' D_{PP}^2 F[q, q] \neq 0$, two equilibria merge and disappear, normal form $x^2 - \tau = 0$; $\mathcal{U}_{\text{sing}}$ is locally a codimension-1 hypersurface

Structural models: payoffs $U = U(\theta)$

Setup:

$\theta \in \Theta \subset \mathbb{R}^{d_\theta}$, typically $d_\theta < d_U$; $\mathcal{E}(\theta) = \mathcal{E}(U(\theta))$, $\Theta_{\text{reg}} = U^{-1}(\mathcal{U}_{\text{reg}})$

Theorem Connected open Θ_0 with $U(\Theta_0) \subset \mathcal{U}_{\text{reg}} \Rightarrow$ constant number of global C^2 branches $P_e(\theta)$, uniformly separated

Proposition $\mathcal{E}(\theta)$ uhc with closed graph; one-sided branch limits when $d_\theta = 1$

Generic regularity does not transfer automatically: $U(\Theta)$ is a thin set

Theorem If $d_\theta = d_U$ and $\det D_\theta U \neq 0$: Θ_{sing} is null, all unrestricted results transfer, folds in U are folds in θ

Theorem If $d_\theta < d_U$, under **structural transversality** at the fold $r' D_U F D_\theta U \neq 0$: Θ_{sing} is locally a codimension-1 submanifold; isolated points if $d_\theta = 1$

Corollary $d_\theta = 1$, compact Θ , folds as above: **finitely many singular points**, Θ splits into finitely many regular intervals, equilibrium graph = finitely many maximal branches

Likelihood correspondence

Definitions:

- ▶ Likelihood correspondence $\theta \mapsto \{\ell(P) : P \in \mathcal{E}(\theta)\}$, branch likelihoods $\ell_e(\theta) = \ell(P_e(\theta))$
- ▶ Full solution envelope (NRLS objective)
 $\ell^*(\theta) = \max_{P \in \mathcal{E}(\theta)} \ell(P) = \max_e \ell_e(\theta)$

Properties:

Theorem Along regular branches $\ell_e(\theta)$ is C^2

Theorem At a scalar structural fold $\theta^\dagger: P_\pm(\theta) \rightarrow P^\dagger$ and $\ell_\pm(\theta) \rightarrow \ell(P^\dagger)$, so branch likelihoods **meet continuously, but with infinite slope**

Proposition If all other equilibria are strictly worse near the fold, $\limsup_{\theta \uparrow \theta^\dagger} \max_{P \in \mathcal{E}(\theta)} \ell(P) < \ell(P^\dagger)$, then ℓ^* **jumps at $\theta^\dagger$** : the profile likelihood is discontinuous

$\Rightarrow$ Analyze the MLE as a constrained problem over the equilibrium graph $\mathcal{G}^* = \{(\theta, P) : F(P; U(\theta)) = 0\}$, not through ℓ^*

Full solution MLE: asymptotics

Estimator: $N \rightarrow \infty$ markets, T fixed, $(\hat{\theta}, \hat{P}) = \arg \max_{(\theta, P) \in \mathcal{G}^*} \mathcal{L}_{NT}(\theta, P)$

Assumption 2 $\Theta, \mathcal{P}$ compact, F continuous; Θ_{reg} = finite union of regimes, each with finitely many continuous branches; singular solutions are limits of regular branches

Assumption 3 (θ_0, P_0) unique maximizer of the population objective on $\mathcal{G}^*$; ℓ continuous with integrable envelope

Theorem Consistency: $(\hat{\theta}, \hat{P}) \rightarrow_p (\theta_0, P_0)$ despite multiple equilibria and discontinuous ℓ^*

In progress:

- ▶ $\sqrt{N}$ -asymptotic normality for regular θ_0 (interior, regular DGP equilibrium, nonsingular information): standard, derived informally
- ▶ Singular θ_0 : nonstandard limit (cf. Andrews 1999)
- ▶ Game primitives deliver the equilibrium-side Assumption 2

CCP estimators: regular (θ_0, P_0) gives the smooth local branch $P_0(\theta)$ assumed in the CCP / pseudo-likelihood literature