The Gorman Lectures.

Structural Analysis of Market Counterfactuals

Ariel Pakes
Harvard University.

May 25, 2025

Terrence's impact on structural empirical work.

,

• Demand systems are at the heart of much of structural empirical work. Terrence did foundational work that formed the basis for much of modern demand analysis. The question he focused on was

How can we estimate demand for a large number of goods with a small number of parameters?

• Terrence's polar forms were necessary and sufficient for multi-stage budgeting; the heart of the Almost Ideal Demand System which was

- used intensively by public finance, labor, macro,....; any situation where one was dealing with broad aggregates.
- but not much in the differentiated products environments integral to Industrial Organization; the polar forms did not allow for sufficient differentiation in cross price elasticities.

• I talked to Terrence several times, but unfortunately never asked him if this was why he turned to characteristic space; but I like to think so.

• Either way, his working paper for the Iowa State College of Agriculture" in 1956 , " A Possible Procedure for Analyzing Quality Differentials in the Egg Market" (reprinted in *Restud*, 1980) was probably the closest forerunner of BLP (1995).

• Both multi-level demand systems and characteristic based models of market demand are, by now, pretty well understood, so these lectures will focus on other aspects of market analysis.

• Still there is another sense in which these lectures will follow Terence's lead; I will take the theoretical structure that provides the framework for the econometric and empirical work seriously, as did Terrence.

• This will include modifications of theory to address issues in my own empirical work which have bothered me.

The lectures.

• I will begin with some notes on the empirical work which foreshadowed the issues I focus on.

• The focus will then turn to two (of the many) topics that have made dynamic counterfactual analysis difficult.

• What they have in common is an attempt to mitigate the cognitive requirements assumed in our models.

• The first is an issue in modeling demand and relies on adapting recent developments in econometrics.

• The second relies on recent developments in theory, and focuses on analyzing market evolution conditional on demand and costs.

• Other predecessors: In addition to Terrence, the Cowles commission was probably the other forerunner of structural analysis, as the following quote makes evident.

"The method of econometric research aims, essentially, at a conjunction of economic theory and actual measurements, using the theory and technique of statistical inference as a bridge pier."

From: T. Haavelmo, in his 120 page treatise, "The Probability Approach to Econometrics", *Econometrica*, 1944.

• Empirical work in I.O. Industrial Organization took Terrence and the Cowles commission quite seriously, and much of the profession seems to be coming along.

• This was facilitated by

- prior developments in theory (both economic and econometric),
- the increased availability of computer power (both for easing computational problems and for the generation of data),
- and the absence of a credible alternative way of analyzing counterfactual environments.

Lessons from productivity analysis (Olley and Pakes,1996).

• The breakup of A.T.&T and productivity (i.e. price $\times$ quantity divided by an index of costs) in the telecommunications equipment industry.

- Empirically: productivity movements are highly predictive of the growth, decline, & exit of firms.
- Partition changes in productivity into : i) reallocation of quantity to more productive plants and ii) changes in plant average productivity.
- That is if $p_t = \sum_i s_{i,t} p_{i,t}$, then $p_t \equiv \overline{p}_t + \sum_i s_{i,t}(p_{i,t} - \overline{p}_t)$.

• Productivity table.

- Partial deregulation, 1975 (registration and certification program). Results in a shift in output to more productive plants
- Consent decree (1982) to completion of breakup (1984): costs of adjustment. Then allocation improves again.

TABLE 1

Decomposition of Productivity

Year	p_t	$\overline{p}_t$	$\Sigma_i \Delta s_{it} \Delta p_{it}$	$corr(p_t, k_t)$
1975	0.72	0.66	0.06	-0.11
1976	0.77	0.69	0.07	-0.12
1977	0.75	0.72	0.03	-0.09
1978	0.92	0.80	0.12	-0.05
1979	0.95	0.84	0.12	-0.05
1980	1.12	0.84	0.28	-0.02
1981	1.11	0.76	0.35	0.02
1982	1.08	0.77	0.31	-0.01
1983	0.84	0.76	0.08	-0.07
1984	0.90	0.83	0.07	-0.09
1985	0.99	0.72	0.26	0.02
1986	0.92	0.72	0.20	0.03
1987	0.97	0.66	0.32	0.10

Two other productivity facts.

• Prior to deregulation (in 1975) plant level capital was negatively correlated with productivity.

- A.T.& T : monopoly & purchased from their subsidiary, Western Electric, which had little incentive to innovate.
- The breakup: Entry (Ericson, Northern Telecom & Hitachi) and the "baby bells" chose their more modern equipment.

• $\overline{p}_t$ declined rather sharply after the breakup.

- This is a period of rapid technological advance in the industry.
- Recall, this is revenue productivity and entry caused prices to fall.

• Using productivity to understand the impact of events on firm growth & the intraindustry allocation of revenue seemed to work, but productivity was less helpful for understanding welfare.

Why couldn't we dig deeper into welfare?

- There were two problems, and by now they are both widely recognized.
  - We had data on revenue, but not separately on price and quantity, and most of the economy is populated by multi-product firms.
    - In a few instances less aggregated measures of price and/or quantity have become available, and have narrowed the problems.
    - The attempt to go directly to markups? Results are being questioned, but either way it lacks a basis for understanding what generated the markups, and so does not lead to credible counterfactual analysis.
  - The dynamic models that were available were demanding. Even if demand and cost systems are known they
    - require cognitive abilities that the agents making the decisions are unlikely to have,
    - led empirical researchers to simplify in ways that, though perhaps the best available analysis, made it harder to believe the empirical results.

Welfare Needs Demand and Pricing Equations.

• Differentiated products demand systems provided new techniques for handling old problems, and have been fairly effective at analyzing a cross section of choices. Two BLP articles (1995, 2004) kicked this off.

• One indicator of its success is how well it performed in determining price in data on cross sections of retail markets.

- Markup Example: Wollman's demand system for trucks. Price on cost characteristics + markup (instrumented coefficient should = 1).
- The good news is how well the static Nash in prices assumption tends to fit cross-sections of product prices in retail markets (though you need to be careful about market definition).
- Hedonics predicted the importance of own characteristics but had no role for competitor characteristics (the markup term).

- Example Pakes (2021): Wollman's (2018) data/estimates/computations.

	Price	(S.E.)	Price	(S.E.)
Gross Weight	.36	(0.01)	.36	(.003)
Cab-over	.13	(0.01)	.13	(0.01)
Compact front	-.19	(0.04)	-0.21	(0.03)
long cab	-.01	(0.04)	-0.03	(0.03)
Wage	.08	(.003)	0.08	(.003)
$\widehat{Markup}$	.92	(0.31)	1.12	(0.22)
Time dummies?	No	n.r.	Yes	n.r.
R^2	0.86	n.r.	0.94	n.r.

- Time dummies pick up most of what the characteristics do not. Fix characteristics and look at the prices of a given car over time (i.e. look at the effect of "competition" on price), the $R^2 \approx 50\%$.

- **Problem that limits current analysis.** Little is known about appropriate pricing in other settings (e.g.'s bargaining, forward looking prices). Hinders counterfactual analysis; if we change investment we typically expect price changes to follow.

Consumer Dynamics and Panel Data on Choices.

• This data is increasingly available, and generates an ability to analyze both consumer demand dynamics *and firm's response to that*.

• I want to distinguish between two situations.

1. One where the decisions made are closely tied to future outcomes.
    - Prominent examples; investment decisions (firms or individuals).
    - Many advances here. Initial papers of Miller, 1984, Wolpin, 1984, Pakes, 1986, & Rust, 1987 (who developed an oft used framework).
2. Decisions which, because the implications of a bad choice can
    - easily be reversed if needed, &/or where repeated evaluations
    - have cognitive costs that rival implied utility gains

• Models of "inattention" &/or "switching costs" seem relevant here, and they lead to a direct dependence on past choices.

- In these cases we need a model for demand *conditional on both*
  - determinants of preferences and
  - past choices.

- Some of the determinants of preferences are not observed and are relatively stable over time making them correlated with past choices.

- So if we put past choices into the model with no correction for unobserved preferences, its coefficient will pick up their impact.

- Other coefficients will also be biased. E.g. for pricing responses we want to distinguish whether the fact that people stay with their last choice when prices change is
  - because they just like their last choice, or because of inattention.

- If the unobserved preferences do not change much over short time intervals we can get credible effects on price by conditioning on individual and choice specific fixed effects.

Model: Pakes, Porter, Shepard and CalderWang (2024).

Observations (i, t) make choices $d \in \mathcal{D}$, where $\#\mathcal{D} = D \geq 2$.

$$U_{i,d,t} = \underbrace{(-p_{d,i,t} - \kappa_0 \cdot 1\{y_{i,t-1} \neq d\}) \cdot \beta_i \; + \; \lambda_{d,i}}_{\text{Structural Utility } (=SU_{d,i,t})} \; + \; \underbrace{\varepsilon_{d,i,t}}_{\text{Error}}$$

with $\beta_i > 0$. The observed choice is $y_{i,t} = \arg\max_{d \in \mathcal{D}} U_{d,i,t}$. Assume

$$\varepsilon_{d,i,t} \big| \, p_i, y_{i,t-1}, \ldots \beta_i, \lambda_i \sim \varepsilon_{d,i,t} | \beta_i, \lambda_i.$$

• $\Rightarrow$ differences in the probability of a choices between any two periods depends only on the structural part of utility (not on the disturbances).

• There is no serial correlation in the errors so we are loading all the time dependence on the state dependence parameter, κ (opposite of Pakes and Porter, 2024).

Price Changes that Induce Switching.

- Notation $\epsilon_{i,t} \equiv [\epsilon_{i,1,t}, \ldots \epsilon_{i,\mathcal{D},t}]$; $\epsilon_i \equiv [\epsilon_{i,1}, \ldots \epsilon_{i,T}]$; $\lambda_i \equiv [\lambda_{i,1}, \ldots, \lambda_{i,\mathcal{D}}]$.
- Agent i chooses c at time s and d at time t, where $s < t$ and $\{d, c\} \in \mathcal{D}_t \cap \mathcal{D}_s$. Then

$$U_{d,i,t} - U_{c,i,t} \geq 0, \text{ and } U_{c,i,s} - U_{d,i,s} \geq 0 \quad \Rightarrow$$

$$(-[p_{d,t} - p_{c,t}] - [\mathbf{1}\{y_{i,t-1} \neq d\} - \mathbf{1}\{y_{i,t-1} \neq c\}]\kappa)\,\beta_i + \lambda_{d,i} - \lambda_{c,i} + \epsilon_{d,i,t} - \epsilon_{c,i,t} \geq 0$$

$$(-[p_{c,s} - p_{d,s}] - [\mathbf{1}\{y_{i,s-1} \neq c\} - \mathbf{1}\{y_{i,s-1} \neq d\}]\kappa)\,\beta_i + \lambda_{c,i} - \lambda_{d,i} + \epsilon_{c,i,s} - \epsilon_{d,i,s} \geq 0.$$

- Adding the two inequalities cancels the fixed effects, and since $\beta_i > 0$ & ϵ is non parametric we can set $\beta_i \equiv 1$.

Samuelson's (1938) revealed preference inequalities with; disturbances, F.E., and state dependence.

If $\Delta\Delta x_{i,t,s}^{d,c} \equiv [(x_{d,i,t} - x_{c,i,t}) - (x_{d,i,s} - x_{c,i,s})]$, & we add the revealed preference inequalities we get

$$-\Delta\Delta p_{i,t,s}^{d,c} - sw_{i,t,s}^{d,c}\kappa + \Delta\Delta\epsilon_{i,t,s}^{d,c} \geq 0,$$

where

$$sw_{i,t,s}^{d,c} = sw^{d,c}(y_{i,t-1}, y_{i,s-1}) \equiv$$

$$(\mathbf{1}\{y_{i,t-1} \neq d\} - \mathbf{1}\{y_{i,t-1} \neq c\}) - (\mathbf{1}\{y_{i,s-1} \neq d\} - \mathbf{1}\{y_{i,s-1} \neq c\})$$

• $\beta_i > 0$ so we can divide by it and maintain the sign, and the disturbances are nonparametric results, so we can set $\beta_i \equiv 1$.

Formalities: Stationarity & Contrapositives

$$SU^d_{i,t}(\kappa_0) \equiv \Big(-p_{d,t} - \{y_{i,t-1} \neq d\}\kappa_0 \Big) + \lambda_{d,i}.$$

• Fix κ, let $t > s$, & order choices by $SU^d_{i,t}(\kappa) - SU^d_{i,s}(\kappa)$. λ_i & β_i are *time-invariant* $\Rightarrow$ the choice with the largest increase in SU_i between t and s is

$$d_1(y_{i,t-1}, y_{i,s-1}, p_{i,t}, p_{i,s}; \kappa) =$$

$$max_{d\in\mathcal{D}} \left[(-p_{d,i,t} - \{y_{i,t-1} \neq d\}\kappa) - (-p_{d,i,s} - \{y_{i,s-1} \neq d\}\kappa) \right],$$

and the choice with j^{th} largest difference is

$$d_j(y_{i,t-1}, y_{i,s-1}, p_{i,t}, p_{i,s}; \kappa) =$$

$$max_{d_j \notin \{d_1,\ldots d_{j-1}\}} \left[(-p_{d,i,t} - \{y_{i,t-1} \neq d\}\kappa) - (-p_{d,i,s} - \{y_{i,s-1} \neq d\}\kappa) \right].$$

d_2 improves the next most, d_3 is the third most...

Note.

- $\epsilon_{i,t}|\beta_i, \lambda_i \;\sim\; \epsilon_{i,s}|\beta_i, \lambda_i$, so differences in probabilities over time depend only on $SU(\cdot)_t - SU(\cdot)_s$.
- $SU(\cdot)_t - SU(\cdot)_s$ are independent of (β_i, λ_i), and depend only on $(-\Delta\Delta p^{d,c}_{i,t,s} - sw^{d,c}_{i,t,s}\kappa)$.

Lemma: Chose any $D(\kappa_0) \in \{d^0_1, (d^0_1, d^0_2), (d^0_1, d^0_2, d^0_3), \ldots\}$ & $t > s$. Then

$$\min_{d \in D_0}\Big[-p_{d,i,t} - \{y_{i,t-1} \neq d\}\kappa_0 - \Big(-p_{d,i,s} - \{y_{i,s-1} \neq d\}\kappa_0\Big)\Big]$$

$$\geq \max_{c \notin D_0}\Big[-p_{c,i,t} - \{y_{i,t-1} \neq c\}\kappa_0 - \Big(-p_{c,i,s} - \{y_{i,s-1} \neq c\}\kappa_0\Big)\Big]$$

$$\Rightarrow \Pr(y_{i,t} \in D_0|p_i, y_{i,t-1}, y_{i,t-2}, \beta_i, \lambda_i) \;\geq\; \Pr(y_{i,s} \in D_0|p_i, y_{i,s-1}, \beta_i, \lambda_i). \;\square$$

• In words If $\forall d \in D_0$, $SU(d)$ improves more than $SU(c)$ $\forall c \notin D_0$, then i will be more likely to chose a $d \in D_0$ in t than in s.

- Probabilities still depend on λ_i, β_i ("incidental" parameters).
- But for large enough κ_0 we know for each λ_i, β_i

$$Pr(y_{i,t} = d_1^0 | p_i, y_{i,t-1} = d_1^0, y_{i,t-2}, \beta_i, \lambda_i) \geq Pr(y_{i,t-1} = d_1^0 | p_i, y_{i,t-2}, \beta_i, \lambda_i).$$

- Multiply both sides by $\Pr(y_{i,t-1} = d_1^0 | p_i, y_{i,t-2}, \beta_i, \lambda_i)$, integrate over (β_i, λ_i), and use Jensen's inequality to obtain

$$Pr(y_{i,t} = d_1^0, y_{i,t-1} = d_1^0 | p_i, y_{i,t-2}) \equiv$$

$$E_{\beta_i,\lambda_i}[\Pr(y_{i,t} = d_1^0, y_{i,t-1} = d_1^0 | p_i, y_{i,t-2}, \beta_i, \lambda_i)] \geq$$

$$E_{\beta_i,\lambda_i}\left[\Pr(y_{i,t-1} = d_1^0 | p_i, y_{i,t-2}, \beta_i, \lambda_i)\right]^2$$

$$\geq \left(\Pr(y_{i,t-1} = d_1^0 | p_i, y_{i,t-2})\right)^2.$$

Now divide both sides by $Pr(y_{i,t-1} = d_1^0 | p_i, y_{i,t-2}))$

$$P(y_{i,t} = d_1^0 | y_{i,t-1} = d_1^0, p_i, y_{i,t-2}) \geq Pr(y_{i,t-1} = d_1^0 | p_i, y_{i,t-2}).$$

Content of Inequality.

- Whether the probability inequalities hold is a property of the data and does not depend on the κ that underlies the calculation of $\Delta SU^{d}_{i,t,t-1}(\kappa) - \Delta SU^{c}_{i,t,t-1}(\kappa)$.
- Our theorem: it must hold if $\kappa = \kappa_0$ ("positive" result).
- If they do not hold, then $\kappa \neq \kappa_0$; the "contrapositive", i.e. if $p \Rightarrow q$, then $\sim q \Rightarrow \sim p$.
- Since $\Delta SU^{d}_{i,t,t-1}(\kappa) - \Delta SU^{c}_{i,t,t-1}(\kappa)$ is linear in κ if the contrapositive is true for a particular κ, it will be true for:
  (i) all $\tilde{\kappa} > \kappa$ (if $sw \leq 0$) or (ii) for all $\tilde{\kappa} < \kappa$ (if $sw \geq 0$).
- So finding a contrapositive generates either an;
  (i) upper, or a (ii) lower bound to κ_0.

- Extension: let κ depend on x, observable characteristics.

Logic of Non-Parametric Bounds

- **Upper bound**. Take a group of agents who are at c in $t-2$
  - Say the relative price of c rises in $t-1$ and a portion switch out.
  - For each time t price change, there is a κ large enough to insure that the fraction of the original group that are in d at t should be greater than the fraction who were in d in $t-1$.
  - Contrapositive (if $p \Rightarrow q$, then $\sim q \Rightarrow \sim p$): if the fraction of the group who are at d in t is smaller than at $t-1$ then κ must be bounded from above.
- **Lower bound.** A group of agents who are at c in $t-3$
  - chose c at $t-2$.
  - Say the relative price of d goes up and some switch out.
  - then the price of c falls to a lower level then at $t-2$ yet less chose c in t than in $t-2$.

Theorem which generates estimators (generalization of above).

Part a. If $s = t - 1$ and the supposition of the lemma holds for all $y_{i,t-1} \in D_0 \equiv D(\kappa_0)$, then

$$\Pr(y_{i,t} \in D_0 \,|\, p_i, y_{i,t-1} \in D_0, y_{i,t-2}) \;\geq\; \Pr(y_{i,t-1} \in D_0 \,|\, p_i, y_{i,t-2})$$

Part b. If $s < t - 1$, then the inequality in the lemma contains $y_{i,t-1}$, and the derivation conditions only on $y_{i,s-1}$. So we look for choice sets D_1 given D_0 such that the lemma holds for all $y_{i,t-1} \in D_1$. Then

$$\Pr(y_{i,t} \in D_0 \,|\, p_i, y_{i,t-1} \in D_1, y_{i,s} \in D_0, y_{i,s-1}) \;\geq$$

$$\Pr(y_{i,t-1} \in D_1, y_{i,s} \in D_0 \,|\, p_i, y_{i,s-1}) \quad \square$$

Health insurance: is low response to price changes a concern?

- The CommCare Data. ("Romney Care" in MA).
- Insurance for those with income less than $3\times$ the federal poverty line (FPL=$10,380 in 2010).
- Data : 2009-2013 in 5 regions and 5 income groups.
- We want switching costs that are not induced by a change in agent's physical environment, just by prices, and we require the same choice set in the comparison years.
- So we drop all individuals who change regions, or face different plan offerings in the years we compare. Also Network drops Partners mid-sample; we split it into two plans.
- Other than Network we have four firms but one enters and one exits areas mid sample (Celticare, 2010, & Fallon, 2011).
- t indexes choice times and we use annual enrollment period.

• Divide this data into cells defined by the Cartesian product of; a) couple of years, b) region, c) plan availability, d) prior choice, e) income group.

• Analyze the implications of switching behavior for the tradeoff between price and switching costs *within* each cell ($=x_i$)

• A weighted average of the cell inequalities are then aggregated into 312 exogenous grouping criteria (couple of years, income, prior choice).

• We want to minimize the difference between the time periods used in the inequalities, to capitalize on the fixed effects controlling for unobservables.

• So use moments: compare choices in t and $t-1$ conditional on $t-2$ for upper bound; and t and $t-2$ conditional on $t-3$ for lower bounds.

• **Extension:** We want to examine how κ changes with observable characteristics; so $\kappa = \kappa(x)$. Of particular interest is whether κ varies with income and/or health status.

Prices Changes.

Need Picture

Why Do Prices Change?

• Figure 1: average prices (prices vary over income groups). Networks varied, but no major change. All cost sharing and covered services were regulated and unchanging.

• Regulator changes rules in 2012. Prior: those below poverty could join any program at zero cost. After: new entrants below poverty could only chose one of the two plans with the lowest prices (previous enrollees could stay in their plan).

• Firms respond differently to "auction". BMC (largest firm with 50% of enrollees below poverty line in 2011) decides to forego new entrants, increase its price, and "harvests".

• NHP and Fallon also increase prices but Network decreases price, and CeltiCare with lowest price, keeps prices low. Lots of price variance.

Data on Switching in Response to Price Change

Need table with switching in and out

Econometric advances (moment inequalities).

• Use an "empirical Bayes bootstrap" (for events that occur rarely, earlier moment inequality methods can rely on moments which implicitly contain negative probabilities).

- Taken from Khan, Ponomavera, and Tamer, (2019).
- Generate *nsim* pseudo random samples that mimic the data sample.
- Find the parameter values that satisfy the moment inequalities from each sample.
- Form confidence sets for the parameters of interest; that is find sets that contain the parameters of interest $1-\alpha$ fraction of samples.

• So the only issue is how to generate pseudo random samples from a multinomial distribution.

• The pseudo samples are found by taking draws from a Dirichlet distribution with parameters set by actual sample.

• Numerators and denominators are separate draws from Gamma distributions whose parameters are determined by the data.

• Draws are on the sequence of probabilities used in the inequalities.

- The upper bound draws each possible sequence of (y_t, y_{t-1}) conditional on y_{t-2} for each cell. If there are 4 choices there are 16 probabilities simulated.
- For the lower bound we simulate (y_t, y_{t-1}, y_{t-2}) conditional on y_{t-3} (64 probabilities).
- For each draw, divide by its standard error (derived from the multinomial formula for each cell).
- Calculate the (potentially set) estimator of κ that minimizes the negative part of the moment inequalities (we used a Euclidean norm).

Results: Single κ ($\kappa(x) = \kappa_0, \forall x$)

Table: Nonparametrics ID Set and Confidence Intervals for κ

Sample	ID Set (κ)		95% CI		99% CI	
	LB	UB	LB	UB	LB	UB
Min cell size $n \geq 20$	\$150	\$168	\$150	\$270	\$150	\$354
Min Cell Size = 50	\$102	\$186	\$78	\$294	\$78	\$450
Min Cell Size = 100	\$78	\$186	\$78	\$450	\$78	\$450

• We use a minimum cell size of 50 with a 99% confidence set throughout, as at that level we match the larger cell size.

• Credible set is wide, but the point estimates from this data set in prior work is well over \$1000; over twice our upper bound.

Picture of Objective Function.

Health and Income Heterogeneity

- Model. Normalize price for the lower income.

$$SU_{d,i,t} = -p_{d,i,t}\big(1+\gamma\{2FPL < I_{i,t} \leq 3FPL\}\big) - \{y_{i,t-1} \neq d\}\big(\delta + \delta_s\{s_{i,t}=1\}\big) + \lambda_{i,d}$$

Those with income less than FPL were fully subsidized.

- Results allowing for differences in κ by income
  - Lower income people are much more sensitive to price.
  - Credible sets: much smaller when we condition on income.

- Allow also differences by health status. $s_{i,t} = 1$ could make one more sensitive to price (lowers wealth) or more hesitant to switch (change providers). Results
  - Sick people are much more sensitive to price (this is a low income population);
  - Upper bound dramatically lower; lower bound essentially zero.

Table: Estimates of κ; switching cost relative to annual price.

	Identified Set		Credible set	
One κ	[102, 150]		[78,450]	
By income (I)	I=1 ($2-3\times$ FPL)		II=0 ($1-2\times$ FPL)	
	I.d. Set	C.S.	I.d. Set	C.S.
	[408,432]	[408,516]	[240,257]	[240, 368]
By income and sick (S)	S=1 index sicker than median			
	I=1, S=0		I=0,S=0	
	I.d. Set [432,528]	C.S. [408,528]	I.d. Set [293,377]	C.S. [240 377]
	I=1, S=1		I=0,S=1	
	I.d. Set [120,156]	C.S. [0,240]	I.d. Set [75,155]	C.S. [0,172]

Estimates from the CommCare Data without F.E..

Compare to our U.B.'s: I=1,S=0; $528: I=0, S=1, $172.

• Logits with no F.E. and our data. As in prior research the more observables you add the lower the switching cost.

- Base Case: $\kappa = \$1382$. Adding observables finally get to
- Plan dummies $\otimes$ (age-sex + region + network + chronic illness), $\Rightarrow$ 249 parameters, $\kappa = \$1019$.
- Add Shepard (2016) three constructed variables he finds highly significant: value and past use of the hospital network and use interacted with Partners. Highly significant but κ still $990.
- Add random initial conditions to base case and κ just increases $1,530.

• Shepard, 2022; rich logit model but without F.E.: $1,164.

Counterfactuals: Are F.E. needed?

- Smooth out price inversion by BMC. Predict share change relative from logit models with and without our κ (no F.E.'s).
  - Actual prices $p(2011) = \$58.4, \;\; p(2012) = \$97.6, \;\; p(2013) = \$46.2;$
  - Counterfactual prices
    $p(2011) = \$58.4, \;\; p(2012) = \$70.6, \;\; p(2013) = \$73.2.$
- 2011-2012 in counterfactual price fall (\$97.6 to \$70.6)
  - model without without F.E..03-.04 pt. share $\uparrow$.
  - model with F.E. predicts .16-.17 pt. share $\uparrow$.
- 2011-2012-2013, over two periods price rises 37% (\$46 to \$73),
  - model without F.E. no loss in share.
  - model with F.E. predicts .20.-25 pt. loss in share $\downarrow$ (vs. actual 2013)
- Effects of miss-specification can accumulate over time.

What Have We Learned?

- **Health Insurance**

  • Demand models that do not allow for flexible unobserved heterogeneity generate estimates of state dependence that are seriously biased upwards and of price effects that are biased downward. Large enough to have meaningful effects on the response to price changes.

  • Likely a result of individual specific differences in provider preferences, that are not picked up by random effects nor random initial conditions.

- **Methodology**: If we allow for F.E., inattention,
  - Paper shows that for a non-parametric ϵ distribution: if switching probabilities $\leq .5$ finite bounds for κ requires stationarity.
  - Parametric ϵ distribution: there are upper bounds on parameters that are easy to implement for any distribution, but lower bounds seem to require distribution specific inequalities.

Back to Market Dynamics (assuming demand and cost functions are known).

- The initial dynamic frameworks used in I.O. assumed[1]

  1. state variables evolve as a Markov process (tractable and often fit in-sample data well; can be issues with stationarity of the process).
  2. and the equililbrium is some form of Markov Perfection (no agent has an incentive to deviate at any value of the state variables).

- The perfectness assumption required agents
  - to access a large amount of information (all state variables), &
  - compute or learn an unrealistic number of strategies (one for each information set, some of which had never been sampled).

- Mitigating the information requirements by allowing for asymmetric information, accentuated the computational problems.

[1] Maskin and Tirole (1988) theory; Ericson and Pakes (1995) applied framework.

- The cognitive requirements of the standard models led to
  - doubts about the model's ability to mimic behavior, and
  - implementation problems for researchers.

- Also problematic in dynamic single agent problems, but I.O. was concerned with market interactions that differed across participants.

- Two types of responses in the empirical I.O. literature. Both were helpful, but not really adequate for counterfactual analysis.
  - focus on "unilateral" effects. E.g. how individual firm's investments respond to incentives (R&D, advertising . . .); problem is the difference between "unilateral" and equilibrium effects of a policy.
  - use approximations: problem is the environment that was approximated differed from the counterfactual environment [2].

[2] See Benkard et. al. 2008.

- Theorists were well aware of the cognitive problems and looked to either
  - weaker notions of equilibrium that were cognitively less demanding,
  - and/or learning processes that mimic how firms respond to change.
- Examples
  - Fudenberg and Levine (1993); Self-confirming equilibrium (Nash conditions only "on-path").
  - Osborne and Rubinstein (1998); Procedurally rational players (play best response but don't worry about computing policies).
  - Esponda and Pouzo (2016); Berk-Nash equilibrium (improper priors, convergence to something close to euilibrium).

• These papers considered games that had state variables that did not evolve over time. In that context the focus was either on learning about competitors' play or characterizing equilibrium behavior.

• There is a also small empirical literature on learning; similarly focused on situations with repeated (e.g.'s bidding in electricity markets)[3].

• The lack of dynamics in the "payoff relevant" state variables, does not allow us to analyze how investment decisions and policies related to them effect the evolution of markets.

• So though the ideas in the theory papers have obvious appeal, to be used in applied work they needed to be adapted to situations where state variables are allowed to evolve over time.

• Experience Based Equilibrium (Fershtman and Pakes, 2012) incorporates analogous ideas into a model that allows state variables to evolve with the outcomes of investment processes and shows that there is a learning algorithm that converges to an EBE.

• We begin with theory and then go to an empirical example.

[3] See Hortascu and Puller, 2008; Doraszelski et. al., 2018.

EBE: Analytic Framework.

- The analytic framework has two parts
  - a notion of equilibrium, &
  - a learning theory that can get us their.

• Equilibrium strategies are "rest points" to a dynamical system. So we need a Nash condition of some sort. Assume

1. firms perceive that they are doing the best they can, and
2. these perceptions are at least consistent with what they observe.

Experience Based Equilibrium: Two Conditions.

1. Maximizing behavior. Agents chose actions which maximize their perceptions of the EDV of future net cash flow conditional on the information at their disposal.
2. Partial consistency of perceptions.Their perceptions need not be correct but they are consistent with what they observe.
   - If a state is visited repeatedly they can learn the distribution of profits at the state, and optimize against that.

• Different agents can have different information sets (asymmetric information) and the theory does not constrain what they contain.

• This can greatly reduce cognitive requirements, but still has a perfectness assumption (just on a reduced state space).

• **Note.** *Though the theory is silent on what the state variables are, empirical work must find a good approximation to them.*

Formalization of Assumptions.

• Denote the information set of firm i in period t by $J_{i,t}$. $J_{i,t}$ will contain both public (ζ_t) and private ($\omega_{i,t}$) information, so $J_{i,t} = \{\zeta_t, \omega_{i,t}\}$.

• The public information, ζ, may differ across firms, and is used to predict competitor behavior, common demand and cost conditions, ...

• Assume $(J_{1,t}, \ldots J_{n_t,t})$ evolves as a finite state Markov process on $\mathcal{J}$.

• Policies, say $m_{i,t} \in \mathcal{M}$, will be functions of $J_{i,t}$. For simplicity assume $\#\mathcal{M}$ is finite, and that it is a "a capital accumulation" game, i.e. $\forall (m_i, m_{-i}) \in \mathcal{M}^n$, & $\forall\, \omega \in \Omega$

$$P_\omega(\cdot | m_i, m_{-i}, \omega) = P_\omega(\cdot | m_i, \omega),$$

(see Asker et. al. 2020 for the modifications needed for a game that does not satisfy this condition; example repeated auctions).

• A "state" of the system, is $s_t = \{J_{1,t}, \dots, J_{n_t,t}\} \in \mathcal{S}$. Note that the agents does not keep track of all of s_t, only $J_{i,t}$.

• $\#\mathcal{S}$ is finite. $\Rightarrow$ any set of policies will insure that s_t will wander into a recurrent subset of $\mathcal{S}$, say $\mathcal{R} \subset \mathcal{S}$ and after that $s_{t+\tau} \in \mathcal{R}$ forever.

• An agent needs a perception of the EDV net cash flow from different policies in order to chose m at state J_i. This consists of an expected profit

$$\pi^E(m|J_i),$$

and a continuation value

$$W(m|J_i), \;\; \forall m \in \mathcal{M} \;\; \& \;\; \forall J_i \in \mathcal{J}.$$

We now formalize our assumptions; recall that they are

- Each agent choses an action which maximizes its perception of its expected discounted value, and
- For those states that are visited repeatedly (are in $\mathcal{R}$) these perceptions are consistent with observed outcomes.

A. $\forall m \in \mathcal{M}$ & $\forall J_i \in \mathcal{J}$

$$W(m^*|J_i) \geq W(m|J_i) \;\; \&$$

B $\forall J_i \in s \in \mathcal{R}$

$$W(m(J_i)|J_i) = \pi^E(m|J_i) + \beta \sum_{J_i'} W(m^*(J_i')|J_i')p^e(J_i'|J_i),$$

where if $p^e(\cdot)$ is the fraction of periods the event occurs

$$\pi^E(m|J_i) \equiv \sum_{J_{-i}} E[\pi(\cdot)|J_i, J_{-i}]p^e(J_{-i}|J_i),$$

and

$$\left\{ p^e(J_{-i}|J_i) \equiv \frac{p^e(J_{-i}, J_i)}{p^e(J_i)} \right\}_{J_{-i}, J_i},$$

while

$$\left\{ p^e(J_i'|J_i) \equiv \frac{p^e(J_i', J_i)}{p^e(J_i)} \right\}_{J_i', J_i}. \quad \spadesuit$$

An algorithm for learning and/or computation

- The algorithm is a generalization of reinforcement or Q-learning. It is iterative and can be used
  - to mimic an actual learning process (then, in principal, no need to impose an equilibrium assumption),
  - or to compute equilibrium (a direct method of computing equilibrium is reviewed below but it doesn't have a learning interpretation).

- **Iterations** defined by
  - A location, $L^k = (J_1^k, \ldots J_{n(k)}^k) \in \mathcal{S}$: information sets of active agents .
  - Objects in memory (in M^k) are:
    - perceived evaluations, W^k,
    - No. of visits to each point, h^k.

- **Must update** (L^k, W^k, h^k)**.**

Updating.

- **Location or L^k.** Calculate "greedy" policies for each agent

$$m^*_{i,k} = arg \max_{m \in \mathcal{M}} W^k(m|J_{i,k})$$

- Take random draws on outcomes conditional on $m^*_{i,k}$: E.g.; if we invest in "payoff relevant" $\omega_{i,k} \in J_{i,k}$, draw $\omega_{i,k+1}$ conditional on $(\omega_{i,k}, m^*_{i,k})$.
- Use outcomes to update $L^k \to L^{k+1}$.

- **Perceptions or W^k.** Simple case: agent observes m_{-i} and knows the primitives $\{\pi_i(\cdot), p(\omega_{i,t+1}|\omega_{i,t}, m_{i,t})\}$.
- Compute ex poste perception of what its value would have been for any m

$$V^{k+1}(J_{i,k}, m) =$$

$$\pi(\omega_{i,k}, m, m_{-i,k}, d_k) + \max_{\tilde{m} \in M} \beta W^k(\tilde{m}|J_{i,k+1}(m)).$$

• Here $J_i^{k+1}(m)$ is what the $k+1$ information would have been given m and using competitors actual play (it is a simultaneous move game).

• Treat $V^{k+1}(J_{i,k})$ as a random draw from the possible realizations of $W(m|J_{i,k})$, and update W^k as in forming averages (simple stochastic integration; Robbins and Monroe,1956)

$$W^{k+1}(m|J_{i,k}) =$$

$$\frac{1}{h^k(J_{i,k})} V^{k+1}(J_{i,k}, m) + \frac{(h^k(J_{i,k}) - 1)}{h^k(J_{i,k})} W^k(m|J_{i,k}),$$

or

$$W^{k+1}(m|J_{i,k}) - W^k(m|J_{i,k}) =$$

$$\frac{1}{h^k(J_{i,k})}[V^{k+1}(J_{i,k}, m) - W^k(m|J_{i,k})].$$

Properties.

• We can only get to equilibrium at points visited repeatedly (in $\mathcal{R}$), but if at equilibrium valuations (* values) we tend to stay their.

$$E[V^*(J_i, m^*)|W^*] = W^*(m^*|J_i).$$

• Algorithm has no necessary curse of dimensionality.

- States: algorithm eventually wanders into $\mathcal{R}$ and stays their. $\#R$ typically grows linearly (not exponentially) in the number of states.
- Computing continuation values: integration is replaced by averaging two numbers.

• The algorithm does not necessarily converge, but one can test for convergence by simulating policies and checking equilibrium conditions without adding much to the computational burden (see paper).

Caveats.

• Using h_j^k, or the number of times visited, as an updating coefficient satisfies Robins and Monroe's conditions for convergence, but is inefficient as later k are more precise, and should be waited accordingly (it matters).

• The algorithm can take a long time to converge

- for major environmental changes, possibly too long to be an adequate approximation to firm behavior.
- Suggests firms learn from more than their own past outcomes at the same state. (see below).
- Finding a learning model that adequately approximates behavior would be extremely useful as it both would provide a transition path and select out (possibly a distribution of) equilibria.

Multiplicity.

• Multiplicity. MPBE is a special case of EBE and generates multiple equilibria; but EBE has an extra source.

• $\mathcal{R}$ contains both "interior" and "boundary" points.

- Interior points are points that can only transit to other points in $\mathcal{R}$ no matter which (feasible) policy is chosen.
- Boundary points. Points at which there are feasible strategies which can lead outside of $\mathcal{R}$.

• Our conditions only insure that perceptions of outcomes are consistent with the results from actual play at interior points. Perceptions of outcomes for some feasible (but inoptimal) policy at boundary points are not tied down by actual outcomes, and different perceptions can lead to different recurrent classes.

Mitigating Multiplicity

- **Knowledge of initial conditions** helps mitigate the problem.
  - Only certain equilibria are consistent with historical situations.
  - In empirical work: asks what would have happened if we changed the environment at the sample's starting date?

- **Boundary Consistency.** Asker et.al.(2014) provide a stronger notion of, and test for, equilibrium. They insure that the choice made at a boundary point is optimal given the information acquired from prior play (or from experimentation). This insures that play at boundary point is optimal given current perceptions (test in paper).

- **Learning models** As noted, if you are willing to chose one it will select out (possibly a distribution) of counterfactuals.

Adding Information to the updating process

• Sometimes the market environment suggests a learning process that can accesses more information

- Hodgson (2024) studies drilling of oil wells. The likely returns from drilling a particular well may be well approximated by the realized gains from nearby wells.
- If there are multiple disjoint markets , one might use information from each market to update all markets. This is particularly appealing when the same firms are competing in those markets (e.g. Zou, in process)
- E.g. Brancaccio et. al. (2020) and other models of transportation networks often use average of past values of matches. If all the counterfactual does is a one time shock to the system, it is easy to recalculate equilibrium. Issues arise if there is a serially correlated aggregate process that you need to keep track of.

- **Non-stationarity.**
  - The worry about post sample non-stationarity influencing in-sample counterfactuals can often be circumvented by using terminal year actions to substitute for perceptions about the post-sample environment
  - Non-stationarity can be built into the learning process in a straightforward manner if the source of the non-stationarity is known, though the meaning of equilibrium has to change.
  - One frequent case is when terminal conditions are known. Then we might want to bring that into the algorithm as the known limit condition, and change the learning process accordingly.

- I now go to an empirical project, where the state variables are not known **a priori**, so the empirical work has to find them. Sometimes we know certain states must be included for determination of policies.