Industrial Reorganization: Learning about Patient Substitution Patterns from Natural Experiments

Devesh Raval, Ted Rosenbaum and Nathan E. Wilson

Discussion by Kate Ho

November 2015

Overview

This is a neat paper

- Very few papers "test" the fit obtained by structural demand models
- An important issue since
  - These models are widely used by researchers and policy-makers but
  - Some researchers have questioned their accuracy
- Clever idea: natural disasters as exogenous shocks to the hospital choice set
- Provides an opportunity to compare models' predictions for resulting changes in consumer choices to realized changes.

Results

Results in many ways not surprising

- Models allowing for flexible interactions between patient characteristics and unobserved hospital quality perform the best
- Models that include hospital characteristics (interacted with patient attributes) but no hospital fixed effects have the worst fit

Results

Results in many ways not surprising

- Models allowing for flexible interactions between patient characteristics and unobserved hospital quality perform the best
- Models that include hospital characteristics (interacted with patient attributes) but no hospital fixed effects have the worst fit

Obvious question: how "good" is the fit of the typical model?

- Most of the paper compares models to each other rather than providing absolute measures of fit
- Statistics for their combination of models: RMSE on predictions of
  - aggregate shares: 0.7 - 2.2%
  - aggregate diversion ratios: 4 - 12%
  - individual level predictions: 19 - 27%
- Fit seems reasonable (at least at aggregate level)
- Equivalent numbers for (best of the) individual models?

Merger Policy Counterfactuals

Final section: implications for hospital merger policy analysis

- Idea: under simple bargaining models (Capps et al 2003), projected change in WTP from a merger (ΔWTP) is correlated with price effect
- This makes WTP, ΔWTP important tools for policy analysis
- Authors want to consider accuracy of the estimated demand models in terms of variables used for merger analysis
  - They compare predicted ΔWTP from counterfactual mergers across models with different RMSE.

Merger Policy Counterfactuals

Potential issue:

- Agree that WTP is an important object in the bargaining game
- Exactly how it enters will vary across models
- Authors assume insurer objective function linear in WTP (Capps et al 2003)
  - So ΔWTP is the right variable to consider
- But in a model where insurers maximize *profit*, and WTP affects insurer demand ($D(WTP)$), things look different
  - Insurer objective function no longer linear in WTP
  - Key object becomes $\Delta D(WTP)$ not ΔWTP
  - (and model predicts price effects of cross-market hospital mergers) (Dafny, Ho, Lee (2015); Ho and Lee (2015))
- I'd be interested to see cross-model comparison for that object too.

Conlon and Mortimer (2015)

Finally: mention related paper by my colleague Chris Conlon

- 2 methods for merger evaluation and measuring diversion ratios
  - Estimate model of demand, predict own- and cross-price elasticities
  - Experiment to exogenously remove a product, observe the products to which consumers actually switch
- They find significant diversion to remaining products
- And show how best to use experimental data to predict price effects of mergers.

Paper is clearly related (and is already cited) - and also helps justify the form of the experiment in this paper.

Conlon and Mortimer (2015)

Relevant point:

- An experiment that eliminates a good has a different effect from one that makes a good slightly less attractive (e.g. slight price increase)
- Reason: the treated population is different across the 2 experiments
  - Experiment 1: average treatment effect on the treated (ATT) for entire population considering the product
  - Experiment 2: local average treatment effect (LATE) for consumers most likely to switch away as price rises
- These populations have different preferences, so the observed diversions will be different.

Conlon and Mortimer (2015)

Question for this paper:

- What's the "right" experiment to use as a test of estimated models, if we want them to be useful to predict merger effects?
- In most settings: post-merger firm chooses price given predicted demand effect of a small price increase
  - We want to test models' ability to predict effects of small price increase
  - An expt that slightly changes prices may be better test than a disaster.
- This setting is different
  - Insurer-hospital pair bargaining over prices
  - Threat point - when hospital dropped from choice set - is key input
  - Model needs to predict where consumers go when choice set changes
  - So the natural disaster may be exactly the right experiment to use.
- Worth mentioning when comparing predicted ΔWTP to RMSE, ie to "accuracy" of model w.r.t. choice when hospitals are removed - probably the *right* measure of accuracy here.

Other Comments

Obvious issue re hurricanes: do they independently affect demand?

- Cheap comment: authors discuss it; not much more they can do
- But of course it could matter.
- Example: LES of Manhattan, quite close to Bellevue, was flooded for some time after Hurricane Sandy.

Finally: how well should we expect these models to do?

- RMSE measures accuracy re: consumer movement across hospitals
- Which is not what the maximum likelihood algorithm is trying to fit
- MLE fits average market shares
- No surprise that fit is best for average shares, worst for individuals
- Or that adding interactions helped a lot.

Overall, encouraging results for these models!