

# CHRISTOPHER CONLON

## RESEARCH STATEMENT

### OCTOBER 2021

The rise of “market power,” the ability of firms to unilaterally set high prices and earn substantial profits, has been the focus of public debate, several high-profile bills in Congress, and major prospective action by antitrust enforcement agencies. The adverse welfare consequences of market power, after all, form the primary basis for most of our antitrust laws. When firms have too much market power, consumers may no longer have access to their preferred products, because firms restrict supply (through capacity shortages or higher prices).

My research addresses the causes and effects of market power among horizontal competitors and also among agents within the vertical supply chain (investors, manufacturers, wholesalers, and retailers). My work includes understanding issues such as: (1) the competitive effects of common ownership — whether the presence of diversified investors reduces competition among firms; (2) how contracts and taxes affect the role of market power within the vertical supply chain; (3) how we measure substitution among competing products.

My approach to measuring and understanding market power blends theoretical models from game theory with modern econometrics, and involves examining one industry at a time and carefully modeling the institutional, legal, and contractual environment in which firms operate. I also regularly interact and communicate with market participants to better understand how they operate.

A related long-term goal of mine has been to improve our econometric tools to better utilize data when measuring substitution between products, and to provide researchers and policymakers more accurate ways to quantify competition and market power.

As part of that work I’ve developed several new methods. Those include better measures of how consumers substitute among competing products (i.e.: *diversion ratios*), as well as understanding how to measure demand and preferences when not all products are available to consumers. It also means using data to better understand *whether* firms are competing (*measuring conduct*). I have also contributed “public goods” to other researchers. These include a software package, developed with a former student, PyBLP, that makes state-of-the-art tools for estimating demand available to a broad set of researchers, policymakers, and antitrust enforcers.

## COMPETITIVE EFFECTS OF COMMON OWNERSHIP

One branch of my work examines a timely set of regulatory questions involving the relationship between managers and their investors. According to the controversial “common ownership” hypothesis, when an investor holds stakes in competing firms, the managers of commonly-owned firms may act to maximize the firms’ joint profits by relaxing competition and avoiding activities that are “business stealing.”

This literature started with IO economists during the 1980s in studies of “joint ventures,” such as agreements between GM and Toyota. More recently, scholars in finance and corporate governance have extended these models to large institutional investors such as BlackRock and Vanguard, who now own between 6-8% of most large publicly traded companies. These ideas have attracted attention from antitrust authorities and policymakers, including: hearings held by the Federal Trade Commission at NYU in December 2018 (in which I participated), the Australian Parliament in September 2021, and in the European Commission.<sup>1</sup>

Along with coauthors Matt Backus and Michael Sinkinson, I examine this “common ownership hypothesis” from a competitive perspective. In a short paper, **Theory and Measurement of Common Ownership** (Published in *Papers and Proceedings of the American Economic Association*) we show that nearly all theories of common ownership make two fundamental assumptions: (a) that investors hold a diversified portfolio of investments and (b) that managers maximize some weighted average of investor/stakeholder payoffs. The former is often measured from financial filings (SEC Form 13F). The latter—how managers aggregate across investors (and which ones), or essentially the theory of corporate governance—must largely be assumed.

We show how to express this general framework in terms of a “profit weight”  $\kappa_{fg}$ , which tells us how much weight the manager of firm  $f$  gives to a dollar of rival  $g$ ’s profit relative to a dollar of their own. The advantage is that  $\kappa_{fg}$ ’s are interpretable in strategic environments, but don’t require taking a stand on *how firms compete*.<sup>2</sup> The main result in our paper **Common Ownership in America 1980-2017** (published in *American Economic Journal: Microeconomics*), is that for two randomly selected S&P 500 components, managers would value a dollar of rival profits at 20 cents in 1980 and 70 cents in 2017. We provide a decomposition and show that the rise in common ownership does not just reflect the fact that diversified investors like BlackRock and Vanguard got large, but that nearly all investors became more diversified.

If managers internalize the effects of their strategic decisions on shareholder portfolios, this would rep-

---

<sup>1</sup>See the final session of the FTC hearing here <https://www.ftc.gov/news-events/audio-video/audio/ftc-hearing-8-common-ownership-audio>. Annex 5 of the EC ruling here [https://ec.europa.eu/competition/mergers/cases/decisions/m7932\\_13668\\_3.pdf](https://ec.europa.eu/competition/mergers/cases/decisions/m7932_13668_3.pdf) and Australia <https://www.aph.gov.au/commonownership>.

<sup>2</sup>The alternatives include focusing on ad-hoc measures such as overlapping 5% or 10% ownership thresholds, or numbers of directors, or other measures that cannot be interpreted as inputs into a larger game. Part of what we do is show how to translate these measures into profit weights. Other alternatives require assuming that firms compete by simultaneously setting quantities of homogeneous products.

resent a substantial reduction in competition and would be large enough to explain roughly 90% of the rise in markups documented by [De Loecker et al. \(2020\)](#). One interpretation of our findings is that the common ownership hypothesis is “big, if true.”

Perhaps the more important (and controversial) assumption is the notion that managers *internalize the effects that their strategic decisions have on their investors’ portfolios*. We explore this question in [Common Ownership and Competition in the Ready-To-Eat Cereal Industry](#) (revise and resubmit at *Econometrica*). The goal for this project is to use the full empirical industrial organization toolkit to test whether or not pricing behavior is consistent with firms maximizing their own profits, or whether managers price in accordance with the “common ownership hypothesis.” Testing models of conduct is a foundational question in industrial organization; our test builds on a long literature in IO going back at least to [Porter \(1983\)](#); [Bresnahan \(1982, 1987\)](#), and applies some recent nonparametric identification results from [Berry and Haile \(2014\)](#).

The goal is to look at data on prices and quantities and determine whether they were generated from a competitive environment or an alternative less competitive environment (such as a cartel or “common ownership”). Given knowledge of consumer demand, different assumptions on firm conduct imply different markups and marginal costs. The challenge is that there is always a sequence of (unobserved) marginal costs that rationalize any outcome. The basis for testing is thus a set of “excluded instruments” that do not affect (unobserved) marginal costs, but might affect demand. This leaves two practical challenges: (1) any test of conduct implicitly tests the specification for the observed portion of marginal costs; (2) the set of potential instruments for testing is infinite, but the choice of instruments will likely affect the answer.

We develop a testing procedure that is meant to address these challenges. We show that the role of instruments is to *predict the markup difference* between the two models. Choosing a model amounts to measuring the correlation between the unobserved marginal costs and the predicted difference in markups.<sup>3</sup> This allows us to reduce an infinite dimensional testing problem into two prediction exercises: (1) predicting the markup difference; (2) predicting marginal costs and calculating the residual (unobserved) marginal cost. Since we are doing prediction, this allows us to use methods from non-parametric regression or machine learning, and means we don’t need to assume a functional form for marginal costs (linear, exponential, logarithmic). Our test can be easily applied outside of common ownership by researchers or policymakers who are trying to detect collusion or determine whether wholesalers or retailers have the power to set prices.

We apply our test to the market for ready-to-eat (RTE) cereal. Questions of market power, collusion, and alleged price fixing have plagued the industry since the 1970s, and it is among the most highly studied

---

<sup>3</sup>The difference in markups functions like the “optimal instrument” of [Chamberlain \(1987\)](#).

markets in the literature (see (Nevo, 2000a, 2001)). Useful for our test of common ownership, there is extensive variation in the financial ownership of Post Cereal, and Kellogg's cereal has a large undiversified shareholder (the Kellogg Family Foundation). Ultimately, we find that managers appear to maximize their own profits, rather than the joint profits of commonly-owned firms.

Some earlier work on measuring competitive effects of common ownership regressed measures of price on measures of market concentration (as modified for overlapping ownership (Azar et al., 2018)). In **Empirical Studies of the Effects of Common Ownership** (*Brookings Working Paper*), we provide a critical review of that literature targeted at practitioners, legal scholars, and a broader policy audience. We show that many of those research designs can generate spurious positive (or negative) correlations, much like an old Structure-Conduct-Performance (SCP) literature that was largely abandoned by IO economists in the 1980s. We also highlight challenges associated with incorporating common ownership into merger review. Specifically, incorporating common ownership effects in merger analyses would increase the implied pre-merger market power, but reduce the *change* in market power arising from the merger. Moreover, we show that merger effects might depend on the financial structure of the deal (all cash vs. share swap, etc.). This review paper was originally written in cooperation with the regulation section at Brookings, and later expanded in cooperation with policymakers at the FTC.

As part of this line of research, we constructed a large database of financial ownership (SEC 13f filings) from the original source documents; it addresses inconsistencies and gaps with commercially-available data. My co-authors and I have made this database available to other scholars studying corporate governance, asset pricing, and competition.

## MARKET POWER AND TAXES IN VERTICAL RELATIONSHIPS

The relationships between manufacturers, wholesalers, and retailers are often important for understanding *how* market power manifests. While there are many testable theoretical predictions, empirical work has been limited by the fact that business-to-business (B2B) contracts are often closely guarded secrets, and access to data is often an important limiting factor for empirical work on these topics.

Julie Mortimer and I have built relationships within the snack-food vending industry that allowed us to partner with a mid-sized operator in the Chicago area (MarkVend). Not only did MarkVend allow us access to essentially all of the data on the enterprise (including contracts with suppliers), but it allowed us to run a number of field experiments using its vending machines.

The most significant project to come out of that work is **Efficiency and Foreclosure Effects of Vertical Rebates: Empirical Evidence** (forthcoming in *Journal of Political Economy*). In this paper we examine *how* a dominant manufacturer can use rebate contracts with a retailer in order to prevent the retailer from selling a rival product (known as “foreclosing” a rival). This particular kind of rebate contract was the sub-

ject of a European Commission antitrust case involving Intel and AMD, and resulted in a fine of \$1.25 billion —one of the largest antitrust fines of all time. In our study, Mars (the makers of Snickers, M&Ms, Milky Way, 3 Musketeers, etc.) offered retailers a per-unit rebate based on the total quantity of Mars products sold, if the retailer met an overall sales target. (These are sometimes called “all-units-discounts” or “bundled discounts.”) The efficiency-enhancing argument for a contract like this is that it offers the retailer a high-powered incentive to restock the vending machine more often and ensure that Mars products are always available. The anticompetitive argument is that the rebate contract may encourage the retailer to fill shelf space with under-performing Mars brands (*3 Musketeers*) and “foreclose” more popular products from Mars’s rival—in this case Hershey’s *Reese’s Peanut Butter Cups* (the #2 selling chocolate confection in the U.S.).

We find that when Mars reduced its retailer sales target, MarkVend responded by swapping *3 Musketeers* with *Peanut Butter Cups*, while also restocking machines less frequently. The net effect was that overall sales went up, and consumers were better off with the reduced sales target. We also estimate a model of demand, simulate purchases, and construct a model of dynamic optimal restocking. Using the model, we find that the net effect of the rebate was that consumer losses from foreclosure dominated the gains from improved restocking. We also find that the rebates were well-targeted: large enough to induce the retailer to stock Mars products instead of Hershey products, but not so large that Mars wouldn’t be willing to pay for foreclosure. Moreover, we show that even if the rival (Hershey’s) gave up all of their profits, they could not avoid foreclosure, even though their product was more popular. In effect, Mars used the market power from top brands (Snickers and M&M’s) and the rebate contract to *tie* its products together.

In [The Price of Liquor is Too Damn High: The Effects of Post and Hold Pricing](#) with Nirupama Rao, we examine the market for distilled spirits in Connecticut. We exploit a law that required both distiller/manufacturers and wholesaler/distributors to post prices with the state Department of Consumer Protection (DCP). After several years, multiple FOIA requests, a number of technological hurdles, and some lobbying of politicians, we were able to gain access to the entire back catalogue of item-level price postings.

Approximately a dozen states (including Connecticut) have some form of “price posting” or “post and hold” law governing the sale of alcoholic beverages. These laws require that wholesale firms post prices in advance of selling, and offer a “lookback” period where firms can meet but not beat the lowest-price competitor. Firms typically are required to commit to selling at (“hold”) those posted prices for 30 days. In the paper we show that the only strategy to survive iterated weak dominance is for each firm to set the monopoly price, and then match any competitor in the second period. This leads states with post-and-hold laws to have higher prices than states without them.

Many have argued that higher prices on alcoholic beverages aren’t necessarily problematic, if higher prices lead to less drinking and associated harms. We show that when compared to alternative policies such as

volumetric taxes, or taxes on alcohol content, these post-and-hold laws are ineffective at discouraging overall alcohol consumption. By creating wholesaler market power they instead lead to large markups on high-end premium products (e.g. Grey Goose Vodka) but relatively low markups on bottom-shelf products. This encourages consumers to substitute to less-desired brands, reducing consumer welfare, while doing little to curb overall consumption. For a given level of alcohol consumption, ethanol taxes lead to higher consumer welfare, and have the additional benefit that they boost government revenues instead of the profits of wholesale distributors.

The broader importance is that policymakers (and economists) may view restricting competition as a “second best” way to address negative externalities in lieu of corrective taxation. An example might be allowing fossil fuel companies to merge rather than introducing carbon taxes, or allowing Coca-Cola and Pepsi to merge in order to combat obesity. Our work suggests this intuition is specific to markets with homogeneous products, and thus invoking consumption externalities to justify lax antitrust enforcement is likely to backfire.

In related work on the market for distilled spirits, [Discrete Prices and the Incidence and Efficiency of Excise Taxes](#) (Published in *American Economic Journal: Economic Policy*) we look at a commonly-asked policy question: How much do alcohol prices respond to tax increases? Previous work suggested that a \$1.00 tax increase implied a retail price increase of around \$1.60. While technically possible, we typically expect firms to *reduce* markups when costs rise. We showed that firms were responding to a (roughly) \$0.20 per bottle tax by either not adjusting prices at all, or by increasing prices in \$1.00 increments, typically so they could maintain prices that end in 99 cents. As one might expect, these sorts of rigidities are largely a retail phenomenon and don’t arise in wholesale prices. Once we accounted for the discreteness in the price response, we were able to flexibly model how the size of the price increase varied with the size of the tax increase. These kinds of questions are important because policymakers often use these estimates to forecast how much revenue can be raised by increasing alcohol taxes. Around 30 states have raised or proposed raising alcohol taxes since 2008, and many have raised other “sin taxes” as well. Congress is currently considering raising taxes on nicotine and tobacco.

The downside of taxes on “sin goods” like alcohol or tobacco is that they are widely believed to be regressive. This led Nirupama Rao and me to write [Who Pays Sin Taxes? Understanding the Overlapping Burdens of Corrective Taxes](#) (with Stern PhD student Yinan Wang). We examine the joint burden across various sin goods (beer, wine, spirits, cigarettes, and proposed taxes on sugary beverages). While cigarette taxes tend to be regressive, wealthier households spend more on average on alcohol taxes.<sup>4</sup> However, in both cases the correlation between sin taxes paid and income is tiny (6% or less). More important than the overall progressivity or regressivity of these taxes is that these taxes are overwhelmingly paid by a small fraction of households (around 10% of households pay more than 80% of taxes). The high degree of con-

---

<sup>4</sup>Even though most alcohol taxes are volumetric rather than *ad valorem*.

centration creates challenges for policymakers seeking to compensate householders for higher sin taxes by reducing other taxes.

In earlier work, I looked at the effect of a different vertical relationship on market power. Since 2000, the single fastest declining component of the consumer price index (CPI) is televisions. In the early 2000s, prices of LCD panels (the main input into televisions) fell around 4-5% per quarter, while finished television prices fell even more quickly (around 6-8% per quarter). This implied that TV manufacturer margins were (rapidly) declining over time. This was hard to square with the fact that the industry was spending large amounts on advertising as well as research and development. In [A Dynamic Model of Prices and Margins in the LCD TV Industry](#) (Working Paper), I show that the combination of declining input prices and forward-looking consumers made the market increasingly competitive, even though the structure of the market remained essentially unchanged during the first two generations of LCD televisions. An important methodological contribution of the paper is to show that dynamic models with forward-looking consumers aren't appreciably more complicated to estimate than static models using the MPEC method of [Su and Judd \(2012\)](#).

## MEASURING COMPETITION AND SUBSTITUTION

Measures of market power, particularly in merger cases, often depend on “market definition” and a firm’s “market share” within that market. In the 2007 merger of *Whole Foods and Wild Oats*, the antitrust authorities argued that the two firms constituted the only two “premium natural and organic supermarkets,” while the merging parties argued that consumers purchase grocery items at a wide variety of stores, including traditional supermarkets, discount retailers (Wal-Mart), and warehouse clubs (Costco).

In 2010, the FTC and Department of Justice Antitrust Division (DOJ) issued updated guidelines governing mergers between horizontal competitors. One main point of emphasis (building on decades of economic research) was the idea that competition is not binary —some products are more substitutable than others, and mergers among similar products lead to larger increases in market power. This led to an increased focus on a measure of competition known as the “diversion ratio.” The guidelines are largely silent, though, on how to measure diversion ratios from data, instead largely hoping that these might be recovered in firms’ “course of business” documents (i.e., internal emails).

In [Empirical Properties of Diversion Ratios](#) (forthcoming in *RAND Journal of Economics*), Julie Mortimer and I address that gap. We show that diversion ratios have an experimental interpretation: (a) raise the price of good *A*; (b) count up the number of consumers who “leave” *A*; and (c) measure the fraction of leavers who choose a substitute *B*. In practice, researchers and antitrust enforcers may not be able to run this ideal experiment. We may instead have access to survey data on second-choices —for example, the UK Competition and Markets Authority surveyed consumers asking “If this supermarket were to close, where would you shop?” In other cases (such as *Sprint/T-Mobile*), we might observe customers switching

between cellular phone providers when network quality changes. We provide a framework to measure and compare diversion ratios from these different interventions. In short, small price changes, changes in product quality, and product removals (or second-choices) all lead to different groups of “leavers” and measure different *average diversion ratios*, and we show how those measures relate to one another. We illustrate our framework in the context of well-known applications from [Berry et al. \(1999\)](#); [Nevo \(2000b\)](#).

When measuring competition among firms, researchers’ main tool is multi-product demand systems, which can be used to evaluate mergers and estimate the value of new products, or measure the contribution of individual products to seller networks. [Berry et al. \(1995\)](#) (henceforth BLP) provide a flexible random coefficients logit model that accounts for the endogeneity of prices, and is the workhorse empirical model for understanding many markets (including consumer products, airlines, pharmaceuticals, hospitals, and schools). The estimator itself (commonly referred to as “BLP”) can be somewhat complicated, as it is a non-convex optimization problem and requires evaluating a number of integrals numerically.

About a decade ago, several papers suggested that estimates from the BLP approach might not be reliable ([Knittel and Metaxoglou, 2014](#)). This led to some debate as to whether there was a problem with the method itself, or whether individual authors had struggled with implementation. A secondary concern is that researchers may sacrifice flexibility in exchange for speed or stability. A long-term project of mine was to provide a standardized implementation of the BLP method, enshrine best practices, and make the method more reliable and accessible to a broad array of researchers.

This ended up being a massive undertaking, which has led to two projects with Jeff Gortmaker (a former Columbia undergraduate, now getting his PhD at Harvard).<sup>5</sup> The first is a large open-source software package implementing the BLP method PyBLP (available at <https://pyblp.readthedocs.io/en/stable/>), which is the largest open-source software project in Industrial Organization. The package itself is pretty fully featured. Users can: estimate demand and supply for many products with fixed effects and consumer demographics; evaluate mergers; and compute standard errors on all predictions.

Time will tell, but our package appears to be getting adopted (around 300 downloads per month) both by antitrust practitioners as well as academic researchers. I am not allowed to divulge specific details, but can say that antitrust enforcers in multiple countries have used our software in some capacity. Our tutorials are also being incorporated in homework assignments in PhD courses in industrial organization.

In the process we learned a lot about the best way to implement the BLP method, and we documented our findings in [Best Practices for Demand Estimation with PyBLP](#) (published in *RAND Journal of Economics*). Many of these findings were ways to avoid numerical problems, and formalized “folk wisdom” that was known by experienced researchers. However, we also provided some useful new results on the

---

<sup>5</sup>A third project (in process) studies how to optimally combine aggregate data on sales and prices with individual data such as that from consumer surveys.



role of instruments, particularly approximations to “nonlinear optimal instruments” when both supply and demand are estimated simultaneously. Perhaps reassuringly, we found that as long as one had strong instruments and followed the best numerical practices implemented in our package, the BLP method is largely reliable.<sup>6</sup>

Another challenge with estimating consumer demand (or preferences) from observational data is that in many retail environments, products are unavailable around 10% of the time. This was a major practical issue in my work with Julie Mortimer on vending machines, where the fraction of products out of stock is often even higher. The naive approach would be to ignore availability and assume that all products are always available. Unfortunately, this understates demand for popular products that stock out more often, and overstates demand for their closest substitutes. In [Demand Estimation Under Incomplete Product Availability](#) (published in *American Economic Journal: Microeconomics*), we develop a procedure to correct demand estimates for the bias arising from out-of-stock events. We show that it is possible to correctly estimate demand with an initial estimate of the demand system and the starting inventory of the out-of-stock product, even when we don’t know when a stock-out event has occurred.

## FURTHER WORK

Lately my research agenda has been dominated by the work on common ownership, some of which is still in progress. One of those projects, related to common ownership and demand-based asset-pricing, is on hold while my student (Daniel Stackman) is on the job market. I have two related ongoing projects. One, with Jeff Gortmaker, extends our work on PyBLP to incorporate “micro data” (such as from surveys or individual purchases). We consider the optimal way to combine aggregate and micro data, and how to construct additional moment restrictions. In the second project, with Julie Mortimer and Paul Sarkis (a PhD student at Boston College), we apply our earlier results on diversion ratios to construct a simple semiparametric estimator using only aggregate market share and second-choice data. Our motivating example is that researchers or antitrust enforcers might have incomplete data on second choices (e.g., we may observe the set of providers that customers who leave AT&T and T-Mobile switch to, but we do not have data from non-merging parties). We show how to use results from our earlier work to construct a semi-parametric logit model using matrix completion.

In the longer term, I hope to partner with firms and run experiments related to contracting. One current effort involves a new platform for the wholesale “swap market” between auto dealers, and a second looks at contracts between franchisors, franchisees, and employees.

---

<sup>6</sup>Understanding the role of instruments in these multi-product demand systems built on my past technical work in [The Empirical Likelihood MPEC Approach to Demand Estimation](#). The main advantage of empirical likelihood over GMM methods is that it eliminates the bias from “many correlated instruments.” BLP-type estimators often rely on instruments such as the average characteristics of other cars, or the number of other sugary cereals. By eliminating this bias, it addressed several of the concerns raised by [Armstrong \(2016\)](#).

## MY WORK

BACKUS, M., C. CONLON, AND M. SINKINSON (2019): “The Common Ownership Hypothesis: Theory and Evidence,” Brookings Working Paper.

——— (2020a): “Common Ownership Data: Scraped SEC form 13F filings for 1999-2017,” <https://doi.org/10.7910/DVN/ZRH3EU>, Harvard Dataverse, V1.

——— (2020b): “Theory and Measurement of Common Ownership,” *AEA: Papers and Proceedings*, 110.

——— (2021a): “Common Ownership and Competition in the Ready-To-Eat Cereal Industry,” Revise and Re-submit: *Econometrica*.

——— (2021b): “Common Ownership in America: 1980–2017,” *American Economic Journal: Microeconomics*, 13, 273–308.

CONLON, C. (2016): “A Dynamic Model of Prices and Margins in the LCD TV Industry,” Working Paper.

——— (2017): “The Empirical Likelihood MPEC Approach to Demand Estimation,” Working Paper.

CONLON, C. AND J. GORTMAKER (2020): “Best practices for differentiated products demand estimation with PyBLP,” *The RAND Journal of Economics*, 51, 1108–1161.

CONLON, C. T. AND J. H. MORTIMER (2013): “Demand Estimation under Incomplete Product Availability,” *American Economic Journal: Microeconomics*, 5, 1–30.

——— (2021): “Empirical Properties of Diversion Ratios,” *RAND Journal of Economics*, Forthcoming.

——— (Forthcoming): “Efficiency and Foreclosure Effects of Vertical Rebates: Empirical Evidence,” *Journal of Political Economy*.

CONLON, C. T. AND N. L. RAO (2019): “The Price of Liquor is Too Damn High: The Effects of Post and Hold Pricing,” Working Paper.

——— (2020): “Discrete Prices and the Incidence and Efficiency of Excise Taxes,” *American Economic Journal: Economic Policy*, 12, 111–43.

CONLON, C. T., N. L. RAO, AND Y. WANG (2021): “Who Pays Sin Taxes? Understanding the Overlapping Burdens of Corrective Taxes,” Working Paper.

Backus et al. (2020b, 2021b, 2019, 2021a, 2020a), Conlon (2017, 2016); Conlon and Gortmaker (2020), Conlon and Mortimer (2021, 2013, Forthcoming), Conlon and Rao (2020, 2019); Conlon et al. (2021)

## OTHER REFERENCES

- ARMSTRONG, T. B. (2016): “Large Market Asymptotics for Differentiated Product Demand Estimators With Economic Models of Supply,” *Econometrica*, 84, 1961–1980.
- AZAR, J., M. C. SCHMALZ, AND I. TECU (2018): “Anti-Competitive Effects of Common Ownership,” *Journal of Finance*, 73.
- BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): “Automobile Prices in Market Equilibrium,” *Econometrica*, 63, 841–890.
- (1999): “Voluntary Export Restraints on Automobiles: Evaluating a Trade Policy,” *American Economic Review*, 89, 400–430.
- BERRY, S. T. AND P. A. HAILE (2014): “Identification in Differentiated Products Markets Using Market Level Data,” *Econometrica*, 82, 1749–1797.
- BRESNAHAN, T. F. (1982): “The Oligopoly Solution Concept is Identified,” *Economics Letters*, 10, 87 – 92.
- (1987): “Competition and Collusion in the American Automobile Industry: The 1955 Price War,” *The Journal of Industrial Economics*, 35, 457–482.
- CHAMBERLAIN, G. (1987): “Asymptotic Efficiency in Estimation with Conditional Moment Restrictions,” *Journal of Econometrics*, 34, 305–334.
- DE LOECKER, J., J. EECKHOUT, AND G. UNGER (2020): “The Rise of Market Power and the Macroeconomic Implications,” *The Quarterly Journal of Economics*, 135, 561–644.
- KNITTEL, C. R. AND K. METAXOGLU (2014): “Estimation of Random-Coefficient Demand Models: Two Empiricists’ Perspective,” *Review of Economics and Statistics*, 96.
- NEVO, A. (2000a): “Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry,” *RAND Journal of Economics*, 31, 395–421.
- (2000b): “A Practitioner’s Guide to Estimation of Random-Coefficients Logit Models of Demand,” *Journal of Economics & Management Strategy*, 9, 513–548.
- (2001): “Measuring Market Power in the Ready-to-Eat Cereal Industry,” *Econometrica*, 69, 307–342.
- PORTER, R. H. (1983): “A Study of Cartel Stability: The Joint Executive Committee, 1880-1886,” *The Bell Journal of Economics*, 14, 301–314.
- SU, C.-L. AND K. L. JUDD (2012): “Constrained Optimization Approaches to Estimation of Structural Models,” *Econometrica*, 80, 2213–2230.