Markup Estimation using Production and Demand Data: An Application to the US Brewing Industry*

Jan De Loecker
KU Leuven
CEPR

Paul T. Scott
Stern School of Business,
New York University

October 23, 2024

Abstract

Relying on data from the US brewing industry, we implement two distinct approaches to estimate markups, using either production or demand data. We find that the two approaches produce similar estimates of average markups and both indicate an upward trend recently. We then combine the two approaches in two ways. First, we replace instruments in demand estimation with a moment involving a production-based markup estimate, yielding similar results. We then evaluate the common (but controversial) assumption that retail markets operate competitively, finding that demand-based markups recovered with the assumption of competitive retail markets agree with production estimates, as long as downstream costs are accounted for.

Keywords: *Market Power; Markups; Demand system; Production Approach; Retail Conduct; Brewing industry.*

*This paper subsumes our earlier paper *Estimating Market Power. Evidence from the US Brewing Industry* (2016, NBER WP22957), and we would like to thank the editor and referees for valuable comments and suggestions. In addition, we would like to thank John Asker, Tim Bresnahan, Chris Conlon, Pierre Dubois, Phil Haile, Igal Hendel, Gaston Illanes, Nate Miller, Aviv Nevo, Mathias Reynaert, Frank Verboven, Matt Weinberg and participants at various seminars for comments and suggestions. We would like to especially thank Nate Miller and Matt Weinberg for generously sharing their data and code to deploy their demand system. Cameron Birchall and Julian Hildago provided excellent research assistance. This paper uses restricted data that were analyzed at the U.S. Census Bureau Research Data Center in New York City. Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper. Contact authors at jan.deloecker@kuleuven.be and ptscott@stern.nyu.edu.

1 Market power and markup measurement

Understanding the extent to which firms exercise market power is a central topic in industrial organization and antitrust economics. A crucial indicator of market power are price-cost margins, or markups, which play a central role in evaluating the impact of potential changes in regulation and market structure (e.g., through a merger).[1]

In order to evaluate market power and changes in market structure, tools were developed to recover estimates of markups based on the premise that the marginal cost of production is not directly observed in the data. This contrasts with an older, and largely discredited, literature that relies on accounting costs, and the assumption that average costs equal marginal costs.[2] Two modern approaches rely on first-order conditions from optimal firm behavior; we refer to these as the *demand-conduct* and *production* approach.

The *demand-conduct* approach recovers the marginal cost of production, and therefore the associated markup, using demand data on, essentially, prices and quantities for a particular market of interest.[3] Given estimated own- and cross-price elasticities across the goods considered, markups can be recovered from the first-order conditions on pricing after specifying how firms compete (e.g., static Nash-Bertrand competition).[4] Traditionally, this approach relies on market-level data, as in Rosse (1970), Bresnahan (1987) and Berry et al. (1995) while some more recent implementations incorporate consumer-level data (Berry et al., 2004; Grieco et al., 2022).

The *production* approach generates markup estimates using production data on variable input expenditures and revenue, paired with the associated output elasticity. By imposing that producers choose a variable input to minimize production costs (conditional on the dynamically chosen inputs, typically capital in most settings), the markup is obtained from the ratio of the output elasticity to the revenue share of the input (given by the expenditure on the variable input over revenue). The production approach was originally developed with industry-level data mind (Hall (1988)), and it has recently been extended to recover markups at the producer (De Loecker and Warzynski (2012)) and product level (De Loecker et al. (2016)).

Each approach has been applied in isolation to various industries and settings. Recently, markup measurement has become central in studying the role of market power, not only in particular differentiated product markets, but also in labor markets and the macro economy, in assessing the impact of economic integration on welfare, and in questions related to

[1] See Asker and Nocke (2021) for discussion.

[2] See Bresnahan (1987) for a discussion.

[3] See Ackerberg et al. (2007) and Verboven (2012) for an excellent overview.

[4] While the standard approach rests on a pricing equation and estimates of the demand system, under certain conditions conduct is testable using the estimated demand system. We refer to Berry and Haile (2021) for an overview.

economic growth, resource allocation, international trade, and innovation.[5] Researchers aiming to address these questions typically cannot rely on the detailed data required for the demand-conduct approach. The production approach has proven to be an alternative in this context, and the approach has become increasingly prominent in applied work.[6] Given the ongoing controversies and research regarding the underlying causes of these trends, obtaining credible measures of markups is a first-order issue.[7]

While both approaches deliver estimates of markups and marginal costs, they have not been applied jointly or contrasted to each other. This paper fills this gap, and we apply the two approaches to the US beer industry. The paper's main contribution is to demonstrate the benefits of deploying the two approaches together: the combination of the two can serve as a validation exercise as well as allowing researchers to test both the economic and econometric assumptions that underpin the approaches.

First, we compare markup estimates from the two approaches using data on the US brewing industry. We find that both approaches provide similar and plausible estimates of average markups. Furthermore, by deploying both approaches across a wide range of data sets, we obtain novel and rich results on the industry-wide time-series of markups. Importantly, both approaches paint a very similar picture and indicate increasing industry-wide markups over time (from about 1.3 in 1972 to over 2 more recently). This result is, at face value, consistent with previously reported (nation-wide) increased levels of concentration (both in production and advertising expenditures) in the US beer industry.[8]

Second, we introduce markups from the production approach as an additional moment in demand estimation, replacing the standard instruments for identifying price sensitivity, cost-shifters. We find that the markup moment serves as a viable substitute for cost shifters, resulting in a similar demand system. These results point to a potential path forward for researchers estimating demand when appropriate instruments are not readily available. This exercise can also be seen as combining the demand and production approaches to test an *econometric* assumption: the validity of the cost-shifter instruments.

Third, we combine the two approaches to shed light on the form of competition in the retail market, illustrating how a joint (production and demand-conduct) approach allows one to verify otherwise ad hoc assumptions about market structure and firm conduct. We introduce a simple and stylized model of retail competition that allows us to test the level

[5] Arkolakis et al. (2015), De Loecker et al. (2020), Peeters (2014), Hsieh and Klenow (2009) and De Loecker et al. (2016).

[6] While the two approaches can be seen as substitutes for some purposes, it should be noted that the production approach cannot be used on its own to address questions that depend explicitly on cross-price elasticities, such as merger simulations.

[7] See Berry et al. (2019) and Syverson (2019) for an overview of some of the most pressing issues and challenges.

[8] See Tremblay et al. (2005).

of competitiveness in the retail sector by combining the approaches. This exercise provides an example of combining the demand and production approaches to test an *economic assumption*.[9]

To connect the demand and production approaches properly, we must consider the industry's vertical structure. While the production approach provides direct information about producers' markups (or marginal cost), demand estimation typically only provides direct information about demand elasticities in final product markets. In general, demand elasticities in final product markets are not equal to demand elasticities faced by producers, as the vertical structure introduces a gap between the prices paid by consumers and prices received by producers, and pass-through of wholesale prices to retail prices may be imperfect. To account for this gap, we introduce and apply a framework that provides an explicit mapping from the estimated product market demand system and extent of retail competition to the brewer's optimal pricing decision.

Our estimates indicate that retail competition is described well by passive retailers that fully pass through wholesale price changes to retail prices, while facing non-trivial retail costs (that are equivalent to markups which are fixed in dollar value terms), as would be the case in a perfectly competitive retail environment. This is at first reassuring news for demand-based studies of producer-level markups and marginal cost that treat the retail sector as non-existent, perfectly competitive, or otherwise imposing perfect wholesale-retail pass-through, from Berry et al. (1995) and Nevo (2001) to Döpper et al. (2021) and Grieco et al. (2024).

However, our results also indicate that the presence of downstream costs (including distribution and retail costs) are important in accurately recovering producer-level markups, and therefore in establishing facts about markups across producers and time. In light of recent work deploying the demand-conduct approach to infer markup patterns over time for producers, researchers may want to take into account potential changes in the retail landscape.

In general, our results indicate that combining production and demand data, and the associated methods, offers an opportunity to study market power along the vertical chain, potentially freeing up some of the hard-to-test assumptions adopted thus far in the literature (see Lee et al. (2021)).

The paper is organized as follows. Section 2 introduces the datasets and some key facts about the US brewing industry, before we discuss the specific implementation of each approach in the context of the US brewing industry. In Section 3 we describe both the demand and production approaches we take to the data, highlighting the specific features

[9] Another economic assumption that can be tested by the combined approaches is the nature of conduct among firms—i.e., testing for collusion. Bet (2021) does just that for the US airlines industry.

of the industry, and we briefly underscore the different trade-offs when adopting these distinct approaches. Section 4 reports and compares the estimated markups using both approaches, and in doing so we document novel facts on industry-wide markups over time.

In Section 5, we combine the production and demand data to 1) evaluate the benefit of adding production-side markup moments in the estimation of demand, and 2) introduce a joint demand-production approach to illustrate how we can recover information about the degree of competition in the retail market. Section 6 concludes.

2 Data on the US Brewing Industry

Our comparison of markup estimates focuses on the US Brewing industry, an important and interesting industry in its own right. The US Brewing industry went through a rapid consolidation period and has served as a prototypical example of oligopolistic market structure, in the US and abroad. It has been the focus of several studies on market power in the US market, including Asker (2016), Rojas (2008), Hellerstein (2008), Goldberg and Hellerstein (2013), Miller and Weinberg (2017), and also in the context of globalization (Alviarez et al. (2020)).

We construct two overlapping datasets: a long panel of US breweries with detailed data on production and input use (plant-level census data), and store-level demand, including weekly prices and quantities at the product level. In this section, we describe the standard demand and production data sets we use for the US Brewing industry. Further details about all data sets are available in Appendix A.

2.1 Demand data

Our demand-based analysis relies on on two datasets: the IRI Academic Database (Bronnenberg et al., 2008) and Kilts Nielsen scanner data.[10] The IRI and Nielsen datasets provide weekly scanner data on consumer purchases at retail stores. The data record the retail prices and volumes. The IRI data span the 2005-2011 period, whereas the Nielsen scanner data cover 2006-2019.[11] These two sources sample different stores across the U.S markets. Each year, the Nielsen data include information for roughly 5000 stores, a sample 4.2 times greater than that of IRI. [12] We refer to the Appendix A for more details on the data construction.

[10] An earlier draft relied on data from Dominick's Finer Foods, a Chicago-area supermarket chain. While this draft focuses on the national IRI and Nielsen data, we report some results based on the Dominick's data in Appendix C.3 .

[11] For more details on the Kilts Nielsen data, see `https://www.chicagobooth.edu/research/kilts/datasets/nielseniq-nielsen`.

[12] These numbers correspond only to stores that sell at least one unit of beer. Stores refer to supermarkets in the IRI Academic Database and food stores in the Nielsen retail scanner data.

2.2 Production data

The data on production comes from the U.S. Census Bureau's Research Data Center Program. We use data on beer brewers, NAICS code 312120, as part of the Census of Manufacturers, and the Annual Survey of Manufacturers from 1972 to 2007. This generates a panel of plants active during the period 1972-2007, and we observe standard production data covering: output, input use (labor, capital, intermediate inputs and energy), investment, and indicators of plant survival (entry/exit).[13] We refer to Appendix A for more details.

While our production data covers a longer and earlier period than our demand data, the data sets overlap in 2007, and we can compare estimated markups using each approach for this year. We also compare the broad time-series patterns of markups using census, Compustat, and the demand datasets, and this way we can inspect the markup time series across the two distinct approaches.[14]

2.3 Limitations and challenges

The data sources we rely on for the demand-conduct approach have been used extensively in the literature, for a variety of research questions including vertical relationships between manufacturers and retailers, pass-through and more generally market power.[15] It is well known that these demand data sources do not cover all relevant consumption, omitting some major retail chains. This could introduce a potential selection in the type of transactions (type of products, more or less price sensitive consumers) that ultimately may impact the demand patterns. Second, the demand data considered in the literature by construction does not capture sales away from home. In the context of alcoholic beverages in the US, the share of this type of consumption is around fifty percent, and stable, during our sample period.[16] The production data, in contrast, will capture all domestically produced beer, regardless where final consumption takes place. In light of our interest in quantifying industry-level markups over time, we alert readers to these limitations of the data.

The production data has its own limitations. In particular, the database constructed

[13]It appears that Census in fact substantially overstates the actual number of domestic brewing companies in existence since it includes entities having establishments primarily engaged in manufacturing any kind of malt beverages, even those that are kept in minimal operation or that are experimental in nature. The Brewing Industry Survey records the active number of breweries – see http://www.brewersassociation.org/statistics/number-of-breweries/. These statistics (and our production data) leave out establishments with a primary classification as something other than a brewery, such as brewpubs.

[14]We refer to De Loecker et al. (2020) for the description of the Compustat database, and how it relates to standard production data (in particular the US census data we use). In particular we select US brewers active in NAICS code 312120.

[15]During the revisions of this paper, a set of papers came out studying markups over time using the Kilts Nielsen scanner data, see e.g., Döpper et al. (2021).

[16]See The Brewers Almanac 2021 report provided by the US beer institute, https://www.beerinstitute.org/.

from the US census of brewers will not include direct imports, and includes exports.[17] The demand data in contrast captures foreign beers that are sold and consumed in the US. The Compustat sample, by design, only captures publicly-traded companies and does not necessarily limit activities to the US market (again including exports), so the reported financial information is thus aggregated across markets.

These are standard challenges that we anticipate to be present in any study that attempts to use one, or both, type of data sources to study industry margins.

2.4 Key industry facts

Figure 1 plots the real value of aggregate shipments, intermediate input expenditures and the total wage bill.

Figure 1: Shipments, Intermediate inputs and Wage bill

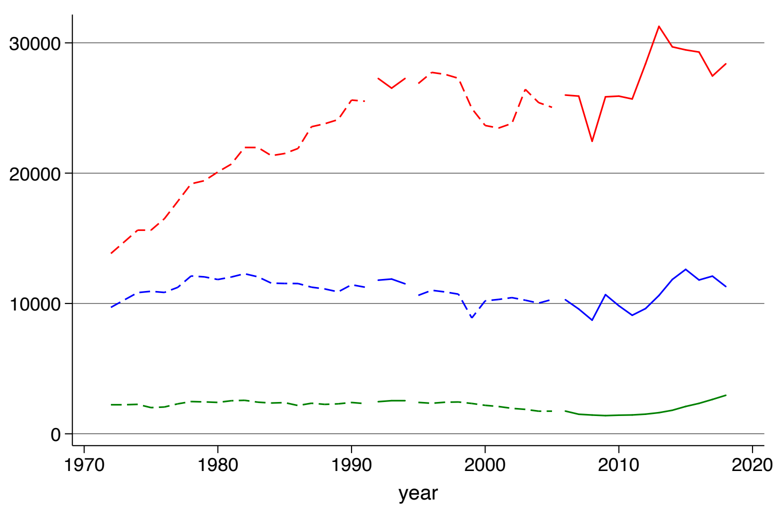

Notes: This figure plots the (deflated) annual value of shipments (in red), intermediate inputs (in blue) and wage bill (in green) as reported by the NBER Manufacturing Database for NAICS 312120. The solid lines indicate for which period we also observe consumption data (i.e. demand data), from Dominick's (1992-1995), IRI (2001-2011) and Nielsen (2006-2019).

Superficially, these trends suggest increasing price-cost margins, as the value of shipments has grown relative to the cost of inputs. Technological change is one potential driver of such changes, but it's also true that the industry has become more concentrated, in terms of sales as well as advertising. Figure 2 plots the nation-wide HHI for production and advertisement over time (for the set of macro brewers). These series indicate an increased

[17] In 2007, a year where we compare the two approaches in detail, US exports and imports were USD 603 million and USD 6.88 billion (out of 138 billion in consumption), respectively.

concentration of economic activity, either measured by sales or advertising, over time in the US beer industry.

Figure 2: Concentration in Production and Advertisement

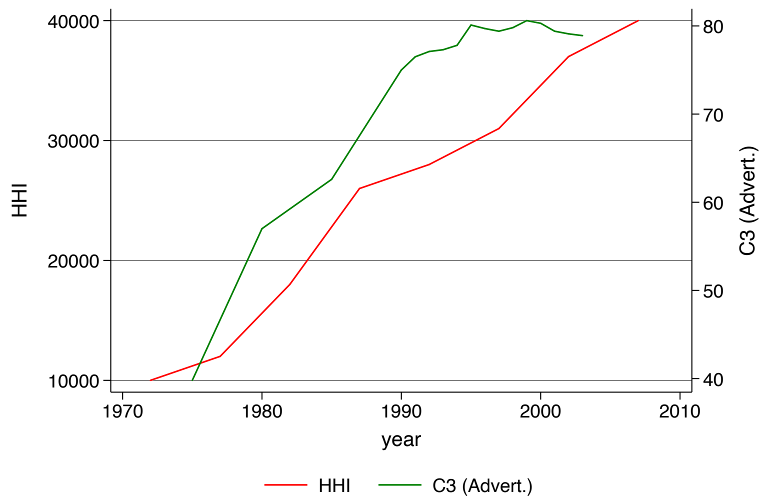

Notes: This figure plots the HHI of US macrobrewers (source: Tremblay et al. (2005) on the right vertical axis. The CR3 ratio for advertising is plotted on right vertical axis (source: Nelson (2005)), as given by the three companies are AB Inbev, SAB Miller, and Coors. Advertising expenditure is measured per case (24 cans).

These trends are, of course, suggestive at most, given the well-known limitations of aggregate concentration measures. That said, these series reflect changes in in aggregate variables and market structure that are consistent with rising markups.

In what follows, we compare markups across the demand and production approaches. We start by comparing markups in a given period of time, before we report the time series of markups. The latter add to an active line of research documenting measures of market power over time (to the extent that markups can be taken as measures of market power).

3 Approaches to markup estimation

The production and demand approaches each require assumptions regarding economic behavior, functional forms, and statistical relationships. However, the approaches rely on distinct sets of assumptions.

Our presentation of each approach is brief as more detailed exposition can be found in many other studies. The demand-based analysis follows the tradition of Rosse (1970),Bresnahan (1987), and Berry et al. (1995) (hereafter BLP), with Conlon and Gortmaker (2020) providing an overview of best practices. Our production-based analysis follows Hall (1988) and De Loecker and Warzynski (2012), hereafter DLW. While De Loecker (2011) outlines the major distinctions between the two approaches, at a conceptual level, very little is known whether the assumptions in either approach are reasonable approximations in practice; and importantly how the two approaches relate to each other, let alone how they can be combined, and what we can learn by combining them.

The demand approach must define a market, specifying the firms, and their associated products, that compete with each other. In addition, in any empirical implementation of a BLP-style demand estimation, an assumption regarding the size of the market is required. The production approach on the other hand groups producers, and their products, to the extent that they produce using a similar production technology.

Because the production approach studies upstream production, and the demand approach considers downstream transactions, modeling of the vertical chain is generally required to make the two approaches comparable (except in cases where manufacturers sell directly to consumers).

3.1 Demand-conduct approach

When implementing the demand approach to recover markups, researchers require 1) estimates of demand elasticities, and 2) an explicit model how firms in the market compete (e.g., static Nash Bertrand). The notion of market equilibrium implies first-order conditions, which together with the demand elasticities, allows one to infer marginal costs and markups (Rosse, 1970; Bresnahan, 1987; Berry et al., 1995). We discuss the two distinct modelling steps, specifying demand for beer and the model of competition.

3.1.1 Demand system

We consider the exact same demand system as Miller and Weinberg (2017), MW hereafter, and when we compute markups across the three demand data sets (IRI, Nielsen and Dominick's) we keep fixed the estimated parameters (estimated using the IRI data). The MW demand system is based on the following model of consumer l in region d at time t, drawing utility from consuming product j:

$$u_{ljdt} = x_j \beta^x_{jl} + \alpha_l P^r_{jdt} + \sigma^D_j + \tau^D_t + \xi_{jdt} + \bar{\varepsilon}_{ljdt}, \tag{1}$$

where x_j is a vector of observable product characteristics, P^r_{jdt} is the retail price, and σ^D_j picks up variation in the valuation of unobserved product characteristics across products (i.e., a product fixed effect), while τ^D_t picks up changes in the valuation of the inside good over time (i.e., a year fixed effect). The error consists of a standard unobserved demand shifter (ξ_{jdt}), and a standard stochastic term. MW adopt a nested logit specification with random coefficients (RCNL), in which $\bar{\varepsilon}_{ljdt}$ can be decomposed into a component that is common to all beer products (but not the outside option), and another component that is product-specific and i.i.d. across products.

MW's demand system allows for some preference parameters to be a function of income. In particular, the price sensitivity parameter α_l is a function of income, and in β^x_{jl}, so are the intercept term, and the coefficient on calories.

We estimate MW's demand system using their restrictions, data, and code. The estimated demand system generates the relevant own- and cross-price elasticities needed (together with the observed market share data), to compute markups and marginal costs after imposing a model of price-setting for the industry. MW's estimation strategy does not impose any supply-side first-order conditions for estimation. That is, the demand system is estimated without making specific assumptions about brewer conduct or about the shape of the cost curves.

3.1.2 Manufacturer markups: vertical structure

In the canonical demand-conduct approach, retail prices are assumed to be directly set by producers (e.g., Berry et al. (1995) and even recent studies such as Grieco et al. (2024)). However, in many industries, the vertical structure drives a wedge between the prices received by manufacturers and prices paid by consumers. This wedge can reflect markups and costs associated with both distributors and retailers. When this wedge is non-trivial (as it is in the case of the brewing industry), care is needed to recover the manufacturer's markups using the demand approach (in contrast, the production approach naturally delivers markups at the manufacturer's level). It may be necessary to correct the approach to account for double or triple marginalization (Berto Villas-Boas, 2007; Bonnet and Dubois, 2010), and even without multiple marginalization (e.g., if retail and distribution sectors are perfectly competitive), it may be important to account for downstream costs.

In other words, to recover markups from an estimated demand system, a model of price determination is required to map from the estimated price elasticities to markups. This model not only needs to describe how producers compete, but also how producer prices map to the retail prices faced by consumers. In this section, we consider what might be described as the prevailing model of vertical interactions in the beer industry.

We start by following the standard assumption in the literature, and specifically applied to the beer market, that the wedge between the brewer and the retail price is entirely cost-driven, or put differently that the retail market clears in a perfectly competitive fashion, admittedly with the presence of non-trivial retail costs (which are not distinguishable from fixed retail markups).[18] In this setting, retail prices correspond to brewer prices plus

[18] There are other models of price setting that would rationalize the idea that brewers set prices, such as

the costs of shipping, distribution, and retailing. Our results highlight the importance of accounting for vertical structure even without double marginalization – downstream costs make a difference when estimating producers' markups. Subsequently, in Section 5.2, we relax the assumption of the passive retail model and instead use the approaches together to identify the degree of competitiveness in the retail sector.

Specifically, our baseline model will rule out double marginalization, meaning that neither retailers nor distributors have independent market power. When it comes to retailers, our motivation for ruling out double marginalization is that retailers operate in a highly competitive environment. Thus, our baseline model treats retailers as adding a fixed amount to the wholesale price that corresponds to the cost of retailing. In contrast, regarding distributors, our view is not that distributors have little market power, but that brewers and distributors maintain close long-run relationships and largely avoid double marginalization problems. We will discuss the plausibility of and motivation for these modeling assumptions regarding distributors and retailers in turn.

There is no precedent in the literature for thinking that there is double marginalization between brewers and distributors. To be clear, brewers and distributors are typically not integrated in a literal sense. Most states impose restrictions on self-distribution, and some have laws limiting resale price maintenance. However, as Asker (2016) writes, the practice of resale price maintenance "does not appear to invite legal sanction." Accordingly, Asker's model involves the brewers setting the prices paid by retailers as well as the prices paid by distributors.

The idea that retailers operate in a competitive environment is probably more controversial, and accordingly we will relax the assumption of competitive retail in section 5.2. MW's modeling is consistent with ours; their main approach rules out double marginalization, and their appendix features an analysis that supports the idea that retailers have little or no market power. Goldberg and Hellerstein (2013) document that pass-through of wholesale to retail beer prices is not significantly different from 100% (in the Dominick's data). Looking beyond studies of beer demand, Berto Villas-Boas (2007) tests vertical structure in the yogurt market, finding support for models without double marginalization (though not necessarily corresponding to competitive retail). In the context of the French bottled water market, Bonnet and Dubois (2010) reject models with double marginalization between retailers and manufacturers. On the other hand, Hoch et al. (1994) conduct an experiment varying supermarket prices and find limited price responsiveness; Chevalier (1995) suggests less directly that competition among supermarkets is imperfect by offering evidence that capital market structure can soften competition among retailers.

Throughout, we use superscripts r to indicate retail variables, e.g. P_j^r indicates the

e.g. retail price maintenance.

retail price of product j, while P_j denotes the brewer's price. For the purposes of this section, where retail markets are treated as competitive, the relationship between retail and wholesale price of product j is $P_j^r = P_j + c^r$, where c^r is the cost of retailing (e.g., including the opportunity cost of shelf space), which we assume to be exogenous. Thus, we assume that retailers passively pass wholesale prices through to retail prices. For all j, we set c^r equal to the (quantity-weighted) average difference between retail and wholesale prices (per 144 ounces) in the Dominick's data in 1994.[19]

Downstream costs matter when recovering producer's markups even if there is complete pass-through between producer and retail prices. In the case of a single-product monopolist, the markup is a function of the own-price elasticity of demand, $\mathcal{E} \equiv \frac{dQ}{dP}\frac{P}{Q}$. With complete pass-through, the derivative of demand with respect to retail prices is the same as the derivative with respect to brewer prices $\left(\frac{dQ}{dP} = \frac{dQ}{dP^r}\right)$. However, $P \neq P^r$, so the retail elasticity of demand is not the same as the elasticity of demand that is relevant to a producer when setting markups. Since, $P_r > P$, retail demand is typically more elastic than the demand a producer faces, meaning that ignoring downstream costs will typically lead to overestimated demand elasticities and underestimated markups.

As discussed above, we treat brewer-distributors as integrated, so P_j can be understood as a wholesale price set by the brewer-distributor. Hereafter, we will simply refer to brewer-distributors as "brewers" and we will account for costs of distribution and wholesaling. Let $s_j(\boldsymbol{P}^r)$ represent the market share for product j when the vector of retail prices is $\boldsymbol{P}^r$ For a given market, brewer b sets prices to maximize their profits, taking the prices of competing firms as given:

$$\max_{\boldsymbol{P}_b} \sum_{j \in J_b} s_j\left(P_j + c^r, \boldsymbol{P}_{-j} + \boldsymbol{c}^r\right)\left(P_j - c_j - c_b^w - \tau\right), \tag{2}$$

where J_b is the set of products produced by brewer b, $\boldsymbol{P}_b$ denotes the vector of prices for $j \in J_b$, and $\boldsymbol{P}_{-j}$ is the vector of brewer prices for products other than j. $\boldsymbol{c}^r$ is a $(J-1)$-dimensional vector with all elements equal to the imputed retail cost c^r. The marginal cost of production is given by c_j.[20] c_b^w represents costs of distribution and wholesaling, which is specific to each brewer and market (explained below). τ represents excise taxes, which are specific to each US state. The product market share function $s_j(\cdot)$ comes from MW's demand system.

[19]We adjust for inflation so that while see an average gap between retail and wholesale prices of 25 cents per 12 pack in 1994, this becomes 35 cents in 2010, which the base year Miller and Weinberg use for inflation adjustments.

[20]To simplify notation, we express marginal cost here as constant, but nothing about our procedure actually requires constant marginal cost. An estimated demand system implies marginal revenue for each product; the Nash-Bertrand first-order conditions simply equate marginal cost with marginal revenue (at the observed prices and quantities), without any need for specific shape restrictions on the cost function.

The first-order conditions for equation (2) can be expressed as follows:

$$\boldsymbol{c}_b = \Delta_b^{-1}\boldsymbol{s}_b + \boldsymbol{P}_b - \boldsymbol{c}_b^w - \boldsymbol{\tau}, \tag{3}$$

where $\boldsymbol{s}_b$, P_b, $\boldsymbol{c}_b$, and $\boldsymbol{\tau}$ are all $|J_b| \times 1$ vectors with rows corresponding to brewer b's products. For a given brewer and market, $\boldsymbol{c}_b^w$ and $\boldsymbol{\tau}$ are constant vectors, with elements equal to c_b^w and τ, respectively. Δ_b is a $|J_b| \times |J_b|$ matrix which is the Jacobian of s_b with respect to P_b (or P_b^r). While we have suppressed market subscripts to simpify notation, we apply this calculation separately for each US brewer in each market (region-month). The main econometric challenge in the demand-conduct approach is obtaining Δ_b.

With our production-based approach, we recover markups for brewers that represent the ratio of a brewer's price to its marginal cost of production, which does not include distribution costs. We want to compute a comparable markup using the demand approach, which is why equation (3) places wholesale costs on the right hand side. Effectively, the demand based approach described above relates marginal revenue to brewer-distributor marginal costs, and if we then subtract the portion of those costs that are associated with the distributor, we isolate the brewer's marginal cost. We must subtract taxes out of the marginal cost in the same way, for they are not included in the markups recovered in the production approach.[21] Distributor costs correspond to shipping costs, which are composed of diesel prices and driver wages. We assume that a truck can transport 900 cases (1800 twelve-packs) and gets 6 miles per gallon. Drivers are assumed to receive \$25/hour and travel at 60 miles per hour.[22] Distribution costs per twelve pack for a product from brewer b in market d and period t are then given by

$$c_b^w = (pdiesel_t \cdot distance_{bdt}/6 + distance_{bdt} \cdot 25/60)/1800. \tag{4}$$

In our final expression for demand-based markups, we also subtract costs incurred by distributors from the price received by brewer-distributors, recovering the price received by brewers.[23] Finally, product j's markup is given by

$$\mu_j = \frac{P_j^r - c^r - c_b^w - \tau}{c_j}, \tag{5}$$

[21]The notion of marginal cost in the production approach is associated with the increased expenditure on a particular variable input, which ignores the cost of excise taxes. However, the marginal cost recovered by the demand approach corresponds to a comprehensive notion of marginal cost which includes taxes. Thus, to make the two approaches comparable, we must subtract taxes as described in equation (3). We obtained federal and state tax rates from the Tax Policy Center: `https://www.taxpolicycenter.org/statistics/state-alcohol-excise-tax-rates`.

[22]See MW for the measures of distances and diesel prices.

[23]We subtract taxes as well as distribution costs from the brewer-distributor's price, reflecting the fact that we use production data capturing the value of shipments that is exclusive of excise taxes and freight charges. Thus, in the context of our production analysis, excise taxes are effectively levied downstream of producers.

noting that everything in the numerator is either directly observed (prices, taxes) or imputed (retail and whole costs), and marginal cost in the denominator is recovered using equation (3).

Our imputation of c_b^w could err in the presence of substantial wage heterogeneity, misestimated travel distances or vehicle capacity, and variable costs not related to driver's wages or fuel. Similarly, our imputation of retail costs c^r is vulnerable to real changes in the cost of retailing over time and heterogeneity in the cost of retailing across retailers. If we were to observe both wholesale and retail prices throughout our sample, we could avoid these robustness issues. With that said, the previous literature has largely ignored these costs, and our results highlight the importance of including these downstream costs when computing producer's markups. Figure C1 illustrates the impact of these imputations—if we were to ignore downstream costs (i.e., implicitly assume that they are zero), the implied markups of brewers are substantially lower.

3.2 Production approach

The *production* approach relies on optimal input demand conditions obtained from cost minimization, and it relates the output elasticity of an input to the share of that input's expenditure in total sales and the markup. To implement the production approach, researchers must first select one or more variable inputs. With an estimate of the associated output elasticity, the conditions for cost-minimization with respect to the variable input allow researchers to compute the markup.

We can formally derive markups by considering how a firm b (a brewer, in our application) can minimize costs while attaining production level Q_{bt}. In the current period, the firm takes as given the levels of dynamic inputs and relevant state variables, so we consider only the first order conditions for variable inputs (V), of which they are potentially several, depending on the context.[24]

Consider the associated Lagrangian function for producer b:

$$\mathcal{L}_{bt} = \sum_V P_{bt}^V V_{bt} + P_{bt}^K K_{bt} + \lambda_{bt}(\overline{Q}_{bt} - Q_{bt}(\boldsymbol{V}_{bt}, K_{bt}, \omega_{bt})), \tag{6}$$

where P_{bt}^V and P_{bt}^K denote the input price for variable inputs V and capital (K), respectively; production finally depends on a brewer-time specific productivity shifter (ω_{bt}).[25]

[24] They typically include intermediate inputs, electricity, and labor. If the researcher allows for multiple variable inputs, overidentifying restrictions can be imposed to estimate markups. Similarly, there can be multiple dynamic inputs in production, and here we use the notation K, since capital is a natural candidate. Our notation will assume one dynamic input to simplify exposition, but this is not necessary; sometimes labor is also treated as a dynamic input.

[25] We use the subscript b to denote the producer. In our application we observe plants, and therefore initially recover plant-level markups. V denotes the variable input, in our application we use employment (later l in logs).

Taking the first order condition with respect to a variable input (V), rearranging terms and multiplying both sides by $\frac{V_{bt}}{Q_{bjt}}$, and let the markup be defined as $\mu_{bt} \equiv \frac{P_{bt}}{\lambda_{bt}}$, we obtain the following expression:

$$= \frac{\partial Q_{bt}(V_{bt}, K_{bt}, \omega_{bt})}{\partial V_{bt}} \frac{V_{bt}}{Q_{bt}} = \mu_{bt} \frac{P^V_{bt} V_{bt}}{P_{bt} Q_{bt}}. \tag{7}$$

This expression holds for *any* model of competition, and this does not even assume there to be a common model of competition across the firms in the data. Because the input price does not have a derivative with respect to the input decision, equation 7 implicitly assumes perfect competition in input markets (or at least exogenous input prices). However, a growing literature extends the approach to allow for input market distortions (Dobbelaere and Mairesse, 2013; Morlacco, 2020; Rubens, 2023; Kirov and Traina, 2023).

To calculate producer-level markups using production data, we require estimates of the output elasticities of one variable input of production and data on that input's revenue share, where the latter is readily available in most production data sets.[26] Just as estimating a matrix of own- and cross-price elasticities of demand is the central task in the demand-conduct approach, estimating output elasticities of production is the central challenge in the production approach.

3.2.1 Technology

Our main, and preferred, specification considers a fixed-proportion (i.e. Leontief) production technology for brewing beer, and we compare this to leading specifications in the literature including gross output and value added. We consider an exhaustive list of empirical models to compute markups by varying the assumptions on i) technology, ii) which inputs to production are not subject to adjustment costs and are freely set by the plant's manager and iii) how productivity evolves.[27]

Throughout we use the following notation (in logs). Production takes labor (l), capital (k), and intermediate inputs (m) to produce output (y), while relying on productivity ω_{bt}. All specifications follow

$$y_{bt} = f^T(m_{bt}, l_{bt}, k_{bt}; \boldsymbol{\beta}^T) + \omega_{bt} + \epsilon_{bt}, \tag{8}$$

where ϵ_{bt} captures measurement error in output. For each technology type T , $f^T(.)$, (capturing gross output (GO), fixed-proportion (FP) and value added (VA)) we also consider

[26] Measurement error in output will be taken into account in our approach.

[27] In line with the previous literature, we abstract away from the potential presence of capacity constraints in production. It is for one likely to impact the measurement of the capital stock, and a such it introduces measurement error in capital (see Collard-Wexler and De Loecker (2016) for a discussion). We have, however, no producer-level data on capacity, or capacity utilization, and therefore we abstract away from this source of model misspecification.

the standard Cobb-Douglas and the Translog functional forms.[28] We discuss the three specifications, before turning to the moment conditions that identify the parameter of interest (on the variable input).

Beer production: Leontief We specify a fixed-proportion production technology for brewing beer:

$$Y_{bt} = \min\left[\kappa_{bt} M_{bt}, F\left(L_{bt}, K_{bt}\right) \exp\left(\omega_{bt}\right)\right] \exp\left(\epsilon_{bt}\right), \tag{9}$$

where Y_{bt} is output, κ_{bt} is the inverse of a fixed per-unit materials input requirement, and $F(.)$ can be any production function in labor and capital inputs. The fixed proportion technology is a plausible description of the brewing process, where the Leontief functional form captures the idea that ingredients—e.g., barley, malt and hops—cannot be simply replaced by using more of labor or equipment and machinery (although labor and machinery might be substitutable for each other, depending on the functional form of F). The bulk of brewing is done at large breweries with a relatively homogeneous production process, reflected in a common technology $F(.)$ across labor and capital.

The Leontief specification has a few advantages. The first two are in terms of the underlying technology and the econometric implications. First, it allows for two sources of technological differences across producers that may be relevant for this industry. The mix of intermediate inputs is left to be plant and time specific (therefore allowing for arbitrary factor-biased technological change with respect to m), which allows for greater flexibility in input use patterns relative to more standard specifications.[29] Second, the technology immediately suggests an expression for the unobserved productivity term which can be used to overcome the standard simultaneity bias concern (Marschak and Andrews, 1944). In particular, the fixed proportion rule, $\kappa_{bt} M_{bt} = F\left(L_{bt}, K_{bt}\right) \exp\left(\omega_{bt}\right)$ implies a simple control for productivity: $\omega_{bt} = \ln(\kappa_{bt}) + m_{bt} - \ln(F(L_{bt}, K_{bt}))$. This control does not require us to take any stance on the degree of the specific form of competition in either the product or factor markets. Substituting the fixed proportion rule into the production function would then consist of literally regressing output net of materials $(y_{bt} - m_{bt})$ on $\ln \kappa_{bt}$, in order to purge the measurement error in output (ϵ_{bt}).[30][31]

Finally, by allowing for a brewer-time specific intermediate mix (κ_{bt}), we allow for an explicit source of product differentiation in the production technology. The ratio of mate-

[28]Translog specifications involve second-order terms and interactions of the inputs.

[29]More precisely, when the only source of cross-firm heterogeneity in the production function is the standard Hicks-neutral productivity shifter (ω_{bt}), and when labor and materials are both flexible inputs and firms face similar input prices, there is no way to rationalize systematically different labor/materials input ratios between firms.

[30]Ackerberg et al. (2015) also make this point in their footnote 17.

[31]If labor and capital are quasi-fixed, and materials is a flexible input, then when output prices are sufficiently low relative to the price of materials, it will be better to set $M = 0$ and not produce at all. However, given that our data only includes actively producing breweries, we do not consider this solution.

rials to the bundle of labor and capital is left to be plant-time specific, and so may be the corresponding input prices. This allows for different plants to purchase different qualities of malt or grains, and different varieties of hops or yeast, resulting in differentiated products. This product differentiation can then show up in observed and unobserved product characteristics in the demand model. Of course, even without heterogeneous κ, product differentiation can still arise from fixed costs activities, such as product design or advertising, in the spirit of Sutton (1991).

Other technologies We compare our preferred Leontief specification to two leading specifications in the literature. First, the gross output (GO) production function considers substitution across all three inputs:

$$y_{bt} = f(m_{bt}, l_{bt}, k_{bt}; \boldsymbol{\beta}) + \omega_{bt} + \epsilon_{bt}, \tag{10}$$

where again the function $f(.)$ is not restricted.

Second, the literature often relies on *value added* to measure output. This is created by netting the value of intermediate inputs from output, which has been referred to as a restricted profit value-added specification. Just as in the Leontief case above, materials are dropped out of the estimating equation. However, the dependent variable is $\log{(Y_{bt} - M_{bt})}$. The micro foundation of this approach is not fully developed – that is, unlike the Leontief specification, the econometric strategy here does not have a clear basis in production theory – and at a minimum, in order to compute the output elasticity of labor (for example), a correction needs to be made:

$$\theta^L_{bt} = \widetilde{\theta}^L_{bt} \frac{(Y_{bt} - M_{bt})}{Y_{bt}}, \tag{11}$$

where $\widetilde{\theta}^X_{bt}$ is the elasticity of value added production with respect to labor, and θ^L_{bt} is the overall output elasticity required to compute markups.[32]

Production approach: moments and estimation Constructing moments for estimation begins with recovering productivity shocks (given parameters). Our strategy follows the tradition of the control function approach, which builds on Olley and Pakes (1996), Levinsohn and Petrin (2003), and Ackerberg et al. (2015) . However, in our main specification, a Leontief production technology suggests a purely technological control function, as opposed as deriving the inverse optimal factor demand. The latter is the approach of OP-LP-ACF which requires including all relevant factor demand shifters, as discussed in De Loecker and Warzynski (2012). In particular, the first stage consists of projecting output on intermediate inputs and a brewer-time specific factor (κ_{bt} in our notation). We approximate this by a function in capital, labor, materials, firm-level wages, year dummies, region

[32]In the case of Cobb-Douglas $\theta^L_{bt} = \beta^L$.

dummies, and interactions (D_s).[33] This specification has the advantage of being the exact same first stage of the ACF routine across the different models of technology we consider (gross output, Leontief and value added). In particular, we follow ACF and rely on a first stage specification:

$$y_{bt} = \phi_t(l_{bt}, m_{bt}, k_{bt}, w_{bt}, D_s) + \epsilon_{bt}. \tag{12}$$

A first-stage control function regression allows us to purge the measurement error from the production function. This first stage serves to separate output variation from measurement error and unanticipated shocks to production: $y_{bt} = y^*_{bt} + \epsilon_{bt}$.[34]

Next, given a vector of production function parameters β, we can recover productivity $\omega_{bt}(\beta)$:

$$\omega_{bt}(\beta) = y^*_{bt} - f^T(., \boldsymbol{\beta}), \tag{13}$$

where, depending on which technology we assume, $f^T(., \boldsymbol{\beta})$ captures the deterministic part of the production function. Given estimated productivities, we can estimate a first-order Markov process for ω_{bt}. Given the process governing ω_{bt}'s evolution, we can compute innovations in the productivity process:

$$\nu_{bt}(\boldsymbol{\beta}) = \omega_{bt}(\boldsymbol{\beta}) - \mathbb{E}(\omega_{bt}(\boldsymbol{\beta})|\omega_{bt-1}(\boldsymbol{\beta})). \tag{14}$$

These innovations, $\nu_{bt}(\boldsymbol{\beta})$, are used to form moments and estimate the production function parameters $\boldsymbol{\beta}$ using the method of moments.[35]

The moments are products of the innovations and inputs. For variable inputs, labor in the case of our main specification (Leontief) we assume $\mathbb{E}(\nu_{bt}(\boldsymbol{\beta})l_{ft-1}) = 0$, while for fixed inputs k we assume $\mathbb{E}(\nu_{bt}(\beta)k_{bt}) = 0$. Contemporaneous capital is always included in the moments. As described in the presentation of the results, we rely on lagged labor, and contrast the results to the case where contemporaneous labor is used (potentially ignoring the response of labor choice to productivity shocks). Lagged material expenditure is included in the moments for gross output production specifications; otherwise, materials expenditure is not included in the moments.

[33] More specifically, we use a second-order polynomial in the inputs, and interactions between the first-order input terms and wages. The dummy variables are not interacted. The regions are East (CT, DE, MA, MD, ME, NH, NJ, NY, PA, RI, VA, VT), Midwest (CO, IA, IL, IN, KS , KY, MI, MN, MO, NE, OH, ND, SD, WI, WV), South (AL, AR, FL, GA, LA, MS, NC, OK, SC, TN, TX), and West (AK, AZ, CA, ID, HI, MT, NM, NV, OR, UT, WA, WY).

[34] Throughout all our specifications we allow for unanticipated shocks to output and measurement error, such that $y_{bt} = y^*_{bt} + \epsilon_{bt}$. As proposed by DLW, we correct the markup formula by eliminating measurement error and use $R_{bt}/\exp(\epsilon_{bt})$ as a measure of revenue.

[35] In general, the identification of production functions relies crucially on assumptions regarding the evolution of the productivity shock and the timing of input decisions (Ackerberg et al., 2015). The functional form of the production function and Hicks neutrality of productivity shocks can be relaxed (Ackerberg and Hahn, 2015). An alternative approach relies on direct measurement of the output elasticities in the data using factor shares, and this approach relies on additional assumptions. See De Loecker et al. (2020) and De Loecker and Syverson (2021) for a discussion.

The validity of the moments on the variable inputs has been a topic of interest in the recent literature on production function identification under perfectly competitive product and factor markets. The presence, however, of plant-specific and serially correlated wages, makes lagged labor a valid moment condition as labor choices are tied over time within a plant. However, our setting is different from that discussed in the literature, and we avoid the need for observed wages for two reasons. First, under the Leontief specification, the first stage is governed by a technology-based control function. Second, as shown by Ackerberg and De Loecker (2024), the presence of oligopolistic competition implies that productivity shocks of competitors (that are serially correlated), and therefore input choices of competitors, move around labor choices. This is sufficient to make lagged labor a valid instrument as serially correlated competitors' productivity again ties labor decisions over time for a given plant.[36]

Production approach: recent advances The production approach has recently received much attention, with studies highlighting the importance of the selection of variable input (Raval, 2023), extending the approach to study input market power (Dobbelaere and Mairesse, 2013; Morlacco, 2020; Rubens, 2023; Kirov and Traina, 2023), reexamining production function estimation in the context of imperfect competition (Doraszelski and Jaumandreu, 2021; Ackerberg and De Loecker, 2024), and taking stock of what the approach has and could tell us about changes in market power (Miller, 2024; Shapiro and Yurukoglu, 2024),

3.2.2 Markups under Leontief production

When using the Leontief specification, we can ignore materials when estimating the production function, and therefore free up the typical challenges associated with identifying its coefficient, but materials cannot be ignored when calculating markups. In our derivation of production-based markups above, we considered production functions that were differentiable with respect to a single variable input. Here, we must consider increasing both materials and labor together to have a well-defined marginal product.

Marginal costs, by producer b and time t, can be expressed as follows:

$$\lambda_{Ybt} = \lambda_{Fbt} + P_{bt}^{M}/\kappa_{bt}, \tag{15}$$

[36] Alternatively, we can rely on an AR(1) specification for productivity, hereby eliminating the heterogeneity in the productivity process, and consider joint moments on the productivity shock and the measurement error. However, this would require using l_{bt-2} as an instrument, which we found to be a weak instrument:

$$\mathbb{E}(\nu_{bt} + \epsilon_{bt} - \rho\epsilon_{bt-1})l_{bt-2})$$

where ρ captures the AR(1) persistence parameter.

where λ_F is the marginal cost of an additional unit of F, and λ_Y is the marginal cost of an additional unit of output. Accordingly, markups can be computed as follows:

$$\mu_{bt} = \frac{1}{(\mu_{bt}^F)^{-1} + \alpha_{bt}^M}, \tag{16}$$

where μ_{bt}^F is the markup formula if we treat labor as the variable input and use the original markup formula – i.e., $\mu_{bt}^F = \widetilde{\theta}_{bt}^L \frac{R_{bt}}{(P^L L)_{bt}}$, and $\alpha_{bt}^M = \frac{P_{bt}^M M_{bt}}{R_{bt}}$, is the revenue share of materials expenditure. We stress that the latter is allowed to vary across producers and time.[37]

3.3 Aggregation

As we are primarily interested in computing industry-wide markups, we aggregate micro-level estimates to national averages. In the case of the demand-conduct approach, we obtain markups for each product-market-month, and for the production approach we obtain plant-year measures. In principle, product-level markups can be obtained using the production approach, but it requires observing product-level revenues and recovering the relevant input expenditures, and this lies beyond the scope of this paper. See De Loecker et al. (2016), Valmari (2023), and Orr et al. (2024) for recent studies of multi-product production.

The markups obtained at the plant level can, however, be viewed as a weighted average of product-level markups, where weights are given by the (unobserved) variable input expenditures associated to each product. This interpretation relies on a common production parameter across products within a plant, an assumption that in the case of our setting is very plausible.[38]

4 Markups in the US brewing industry

For each approach, we compute the share-weighted mean markup for US producers for each year in the various samples.[39] We use revenue weights (s_{bt} in the case of plants, and s_{jt} in the case of products) in order to guarantee a meaningful comparison across the two approaches.As large brewers have a large majority of the market, our mean markup

[37] This highlights that we do not require separate info on intermediate input use and prices, while capturing intermediate input heterogeneity across producers and time.

[38] To see this, consider equation (7) for a given product j and let the brewer-level markup be given by $\mu_b = \sum_j a_{bj}\mu_{bj}$. When $\theta_{bj} = \theta_b \; \forall j \in b$ and $a_{bj} = \frac{E_{bj}}{E_b}$, where E_{bj} denotes the appropriate expenditure on the variable input, we obtain the markup expression in equation (7).

[39] For the production approach, it's necessary to focus on US producers as we rely on the US Census of Manufactures. When implementing the demand approach, we can in principle recover markups for both domestic- and foreign-produced beers as both are in the data. To maximize comparability to the production approach, when we compute average markups, we only average over those beers produced within the US. However, the production data includes exports, while the demand data does not.

estimates can be understood as corresponding roughly to mean markup estimates for mass-produced beer.[40]

In this paper, we do not intend to innovate in the estimation of demand, so we rely on MW's well-known model of beer demand.[41] We have taken MW's main specification (which they refer to as "RCNL-1") without modification. Our production-based results were disclosed from the US Census Research Data Center in November, 2016, and MW's replication materials became available subsequently.

4.1 External information

To put our markup estimates from the two approaches into context, we present a reported breakdown of the price of a nationally representative 6-pack of beer (Consumer Reports, 1996). Table A1 in Appendix A decomposes the final final price of a 6-pack into the various cost of production categories (ingredients, packaging, labor in production and advertising/management), distribution and retail. We use this to illustrate the margins along the vertical chain, and in particular to get a sense of the brewer's markup. This report suggests markups that are about 1.52 (in 1996) overall (across mass and craft producers), when we take the view that advertising and management are not part of the marginal cost of production. This breakdown further highlights the wedge between retail and producer-level prices, here a wedge of 1.25.

Alternatively, if we treat advertising and management as non-variable expenditures and other costs as fully variable and proportional to output, the report implies a producer-level markup of 1.44 for mass-produced beers. If we treat labor as a non-variable expenditure, the markup changes to 2.27. We view these as rough upper and lower bounds on the plausible range for the average markup in the industry, and they serve as a reference point to validate our markup estimates.

4.2 Markups using the demand approach

We start by isolating one year in our data, 2007, for which we have nation-wide coverage across US markets in the IRI data and the Census of Manufactures data. Moreover, 2007 is a convenient year for this comparison as it precedes the SABMiller-Coors joint venture. Miller and Weinberg (MW) suggest that the merger may have had coordinated effects, meaning

[40] The share of mass brewers declined in recent decades but remains high, ranging from 97 percent in 1980 to about 82 percent in 2012. If we rely on an alternatively sample restriction where we select major US brands in the three demand datasets (Dominick's, IRI and Nielsen), this does not impact the markup estimates substantially.

[41] While MW's demand system allows for preferences to be a function of income, one could argue for the inclusion of richer demographics and different dimensions of heterogeneity. In the appendix, we include markup estimates based on the demand model of of Hildago (2023), which features richer heterogeneity and demographic information.

the Nash-Bertrand model of conduct (and therefore demand-based markup estimates based on the Nash-Bertrand assumption) may no longer apply for years following 2008.

We first highlight the importance of the conduct assumption as well as the retail-cost correction (while keeping with the assumption of perfect competition in the retail market, as presented in section 3.1.2). Table 1 presents average markups for US brewers obtained from the demand approach using MW's demand system and the 2007 IRI data.

Table 1: Markups: Conduct and Vertical Structure (IRI, 2007)

Brewer competition	(1)	(2)	(3)
Nash Bertrand	2.29	2.15	1.80
Nation-wide monopoly	14.91	14.88	14.68
Product-by-product	1.57	1.52	1.39
Retail cost correction	Yes	No	No
Distribution cost correction	Yes	Yes	No

Notes: We report sales-weighted markups across all markets using the estimated demand system of Miller-Weinberg, and we do this for 3 distinct market structures: 1) Multi-product Nash-Bertrand, 2) collusion among domestic brewers, and 3) Nash-Bertrand with each product owned by a single-product firm. For each market structure, we present results where we evaluate the correction of either distribution and/or retail costs (as explained in the main text). We impute a markup of 15 whenever demand-based markups would otherwise be negative or larger than 15.

A minor difficulty when presenting demand-based markup estimates is dealing with negative values.[42] Loosely speaking, negative values can be viewed as very large markups. To see why, consider a single-product example. When the elasticity of demand approaches unity from above, markups approach infinity—an elasticity of -1.01 gives a markup of roughly 100. While this is a massive markup, the very similar demand elasticity of -0.99 gives a negative markup. To deal with this unfortunate discontinuity, whenever our raw estimate of the market is negative, or greater than 15, we impute a markup of 15. Thus, when the monopoly case presented in Table 1 displays a mean markup close to 15, that actually reflects an estimated distribution of markups that is larger than 15 and frequently negative. We follow this imputation practice in all subsequent results, as it avoids strange discontinuities in comparative statics that can occur if we were to keep negative markups (or drop them).

While our preferred notion of price determination is Nash-Bertrand competition, we also consider two alternative assumptions to illustrate how conduct assumptions matter: 1) an extreme assumption where all North American Brewers collude to set prices as if all their

[42] There does not seem to be a systematic way the literature handles this issue, rather ad hoc solutions are adopted, like dropping those observations, or ex-post adjusting the model specifications.

brands were owned by one firm, and 2) that each product is priced to maximize profits from that product—i.e., we pretend each product is managed by a single-product firm.

Unsurprisingly, the assumption of full collusion leads to (perhaps unreasonably) high markups. Assuming product-by-product pricing leads to somewhat lower markups than Nash-Bertrand.

Another assumption we vary is whether downstream costs are accounted for. Specifically, we separately evaluate the importance of correcting for distribution and retail costs, and relate it to when we completely ignore the industry's vertical structure, imputing zero wholesale costs and imposing that brewer prices equal retail prices (in our notation $P = P^r$) when recovering brewers' costs with equation (3).

Our preferred estimate for the 2007 markup is thus 2.29 (in the IRI data), using a multi-product Nash-Bertrand conduct assumption on brewer price-setting, with a passive retail sector that faces retail and distribution costs. The presence of these costs play, as expected, an important role in the level of the estimated markup; implying a range of $1.8 - 2.29$ depending on whether both downstream costs are taken into account or not.

4.3 Markups using the production approach

Table 2 presents the markup estimates across different production function specifications, labor or materials (when appropriate) as the variable input, and instruments for the innovations in productivity. To facilitate direct comparison with the demand-based estimates, we present revenue-weighted average markups for 2007.

For some of the gross output specifications, we obtain very large standard errors on markups, reflecting the difficulty of identifying the production function parameters for variable inputs, especially in the translog case. Within the gross-output estimates, only the case of Cobb-Douglas technology with materials as the variable input has reasonably precise standard errors.

The Leontief point estimates (2.08) are very close to the point estimates on the demand side, and the standard errors are reasonably small (.06).[43] Looking further back, we estimate production-based markups of 1.83 and 1.89 in 1992 and 1997, respectively (with standard errors $\approx$ .06 in both cases).[44]

We surmise that, at least in the context of modeling breweries, there is much to recommend Leontief specifications. First, the apparent lack of independent identifying variation in

[43] The production function coefficients are presented in Table C2, and one notable observation is that the coefficients on labor are estimated with relatively small standard errors for value added specifications, but labor and materials coefficients both have large standard errors in the gross output specifications, reflecting high correlation between labor and materials inputs.

[44] Our production-based estimates have not changed relative to De Loecker and Scott (2017), but we now focus primarily on the values of markups in 2007, whereas earlier drafts focused on 1992, reflecting a change in overlap with demand data.

Table 2: Markups: Production Technologies (Census, 2007)

Technology	Gross output	Value added	Leontief
		Cobb-Douglas	
Variable M	1.52		
	(0.32)		
Variable L	2.02	5.00	2.08
	(1.78)	(0.86)	(0.06)
		Translog	
Variable M	1.70		
	(0.83)		
Variable L	3.45	1.05	2.05
	(5.23)	(1.09)	(18.10)

Notes: Sales-weighted mean markups across all breweries for 2007. Standard errors (obtained using block bootstrap) in parentheses. For Leontief specifications, Variable-L-based markups include a proportional materials adjustment as described in equation (16). For value-added specifications, the value-added elasticity is multiplied by the ratio of value added to revenue to convert to an output elasticity.

materials and labor inputs makes it difficult to estimate a gross output production function with both materials and labor on the right hand side. This lack of independent identifying variation is also unsurprising if the data were generated by a Leontief production function (equation (9)). Even assuming that all breweries share the same proportionality ($\kappa_b = \kappa$, $\forall b$) between materials and output, there would still be some independent variation in labor and materials due to productivity variation. However, this independent variation will not be identifying variation – given the placement of the error terms in the functional form, materials should be a strictly better predictor of output than labor.[45]

Second, materials plausibly come in a fixed proportion with beer output, as there is little substitutability for beer ingredients. It is worth noting that our estimation strategy allows the proportionality κ_{bt} between materials and output to vary across breweries, or even within a brewery over time, meaning that our approach is compatible with different breweries buying inputs at different prices. This proportionality only enters into our framework when computing markups through the α_{bt}^M term – the ratio of materials expenditure to sales.[46]

[45] We could make this point stronger by placing the productivity term in equation (9) outside the minimum function. In this case, labor and materials will be perfectly collinear, so there is obviously no hope of identifying coefficients on both variables. In the data, labor and materials are not perfectly collinear, so we prefer the form in equation (9), which rationalizes both the *independent variation* in the two variables and lack of *independent identifying variation.*

[46] Heterogeneous input requirements also allow us to capture the fact that some brewers purchase malt directly while others have malting equipment and purchase grains, as long as this does not affect the substitutability of labor and capital. This is the only significant possibility for substitution between materials expenditure and the other inputs that we are aware of.

We convert the reported value of plant-level shipments to a measure of output. We follow the practice of eliminating common price trends over time, and by incorporating state dummies we allow these to vary across regions. Of course, as pointed out by Klette and Griliches (1996) in the context of production function estimation, any unobserved plant-level price variation can potentially introduce a bias in our estimated output elasticity of labor. We are, however, not interested in recovering a measure of physical productivity; and therefore we are only subject to a potential bias in the level of the markup, to the extent that the output elasticity of labor is biased. Klette and Griliches (1996) and De Loecker (2011) find that when correcting for unobserved output prices, the estimated production function coefficients are (roughly) 1.1 times higher (suggesting a negative correlation between the level of inputs and output prices).[47]

We can further rely on our estimates of the demand system to get a more precise sense of the potential bias in the estimated output elasticity of labor. The mean own-price elasticity is about -4.5, suggesting a factor of 4.5/3.5 bias on the output elasticity. To evaluate the impact on markups, we have to revisit equation (16), and this would imply that the markups in 2007, estimated to be 2.076 under the production approach, would become 1.96. This is within the range of a 95% confidence interval for our production estimate (as well as the estimated markups using the demand approach).

4.4 Comparing markups across approaches

We summarize our findings, relying on our preferred demand and production specifications, and report markups across the two approaches in the year 2007, for which we have nation-wide coverage in both data sets.[48] Table 3 summarizes the main findings using the various datasets, and we report markups for the production approach using census data, and also present those using the Compustat sample. We take firm-level markup estimates from De Loecker et al. (2020) for US based firms active in the NAICS code 312120, and compute a sales-weighted industry average by year. The production and demand approaches are thus in broad agreement on the overall level of market power in this industry.

Industry-wide aggregate markups over time A recent literature has emerged that tracks measures of market power over time, in the form of reporting industry-wide aggregate

[47] Under a setup of Klette and Griliches (1996), the estimated output elasticity, when using revenue as a measure of output, is given by $\beta\frac{\eta+1}{\eta}$, with η the elasticity of demand. The bias need not to go this way, as high prices can reflect higher quality, which can materialize through higher input use, conditional on productivity and the same production function across plants. The potential bias only impacts the level of the output elasticity and the level of implied markups, leaving the variation in markups over time unaffected.

[48] We report additional results in the Appendix for an earlier year, 1992, for which we have demand data, albeit only for the Chicago market. The production approach naturally delivers a national average, as we rely on US census data, and the Nielsen and IRI data sets feature national samples of stores.

Table 3: Markup Estimates Across Approaches and Data Sources

Year	External	Production		Demand	
	Cons. Rep.	Census	Compustat	IRI	Nielsen
1996	1.52	-	1.57	-	-
2007		2.08	2.02	2.29	2.18

Notes: We report sales-weighted average markups using our four datasets, and contrast it to the implied markups from the 1996 Consumer Report (presented in Appendix A).

markups.[49] Figure 3 plots industry-wide share-weighted aggregate markups for each each approach and each sample.

The solid red line plots the aggregate markup using the production approach applied to the plant-level census data. This series suggests that overall markups have trended up, going from under 1.5 to just over 2. The same approach applied to the COMPUSTAT sample of US brewers yields a similar time series for aggregate markups. Our micro-level census data stops in the year 2007, leaving a relatively short window to compare the time series of the two approaches. However, there is strong agreement that markups have increased since 2000. For the more recent years, the two different demand samples (IRI and Nielsen) again paint the same picture, a rising aggregate markup.[50]

Markups obtained from either the production or demand approach thus indicate a clear increase in aggregate markups, consistent with the independently reported increase in industry concentration ratios (both in production and in advertising). The fact that the aggregate markup trends are very similar across datasets and approaches is of course a first important result. In addition, the agreement between the results coming from the plant-level census and firm-level Compustat sample is noteworthy, albeit the inherent differences in the scope of the reported items. Equally important is the agreement between the demand-based approach and the production approach applied to Compustat.[51]

To construct the demand-based time series in Figure 3, we assume that MW's parameter estimates are valid outside of the sample used for estimation (2005-2011).[52] This is weaker than assuming that the demand system is fixed over time—preferences are a function of income, and so the demand system changes along with the income distribution. One might

[49]De Loecker et al. (2020) do this for a wide range of industries in the US using COMPUSTAT and census data, Döpper et al. (2021) use the entire Nielsen database, while Grieco et al. (2024) focus on the US car industry.

[50]The latter should be no surprise given that we hold the demand parameters fixed across the demand samples, and differences would only come from the differences in the sample and coverage.

[51]See De Loecker et al. (2020) for a detailed discussion between the trade-offs of using Census and Compustat data.

[52]It is possible that consumer's preferences have changed in ways unrelated to income, which means that the extrapolation of MW's demand system beyond the sample period should be taken with a grain of salt.

Figure 3: Aggregate markups across Methods and Datasets

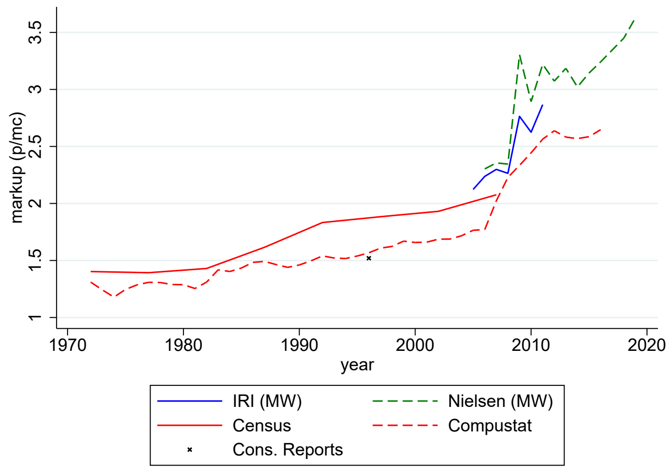

Notes We plot the sales-weighted average markup across firms (using plants, product-market micro estimates of markups) using the production approach (red lines) on the census and Compustat data, and using the demand approach (blue and green lines) on the IRI and Nielsen data. We insert the Consumer Report (1996) based markup for comparison. All demand-based markups are computed assuming the benchmark passive retailing assumption, where we correct for the downstream retail and wholesale costs, as described in the text. Following MW, the merger period of June 2008 to May 2009 is omitted from the demand-based calculations. 2017 was dropped from the Nielsen time series because the sample of stores was very different that year.

argue that the demand system should not be extrapolated in this way—it could be the case that heterogeneity in preferences is correlated with income, but it is not necessarily a direct function of income, and so preferences might not change as the income distribution changes (or they might change less than our approach implies). Therefore, in Appendix C.1, we explore the impact of the changing income distribution, finding that if preferences are held stable over time, we still see the increase in markups associated with the Miller-Coors merger in 2008, but markups are basically flat after that, in contrast to the upward trend we see in Figure 3.

Conduct is of course another important input in the calculation of markups using the demand-conduct approach. In constructing Figure 3, we follow common practice in the literature in assuming Nash-Bertrand competition. Our baseline assumption is that the Miller-Coors merger changed only the ownership matrix and that Nash-Bertrand conduct persisted after the merger. However, MW find evidence that the merger had coordinated effects, and so in Figure C1, we also compute markups using MW's model of conduct, which results in higher markups after 2008.

4.5 Markups, prices, and implied marginal costs

Markups can increase when either prices increase or marginal costs decrease. Figure 4 illustrates that the increases in the above demand-based markups time series are coming from declines in marginal cost that do not fully pass through to prices.

Figure 4: Prices and estimated marginal cost

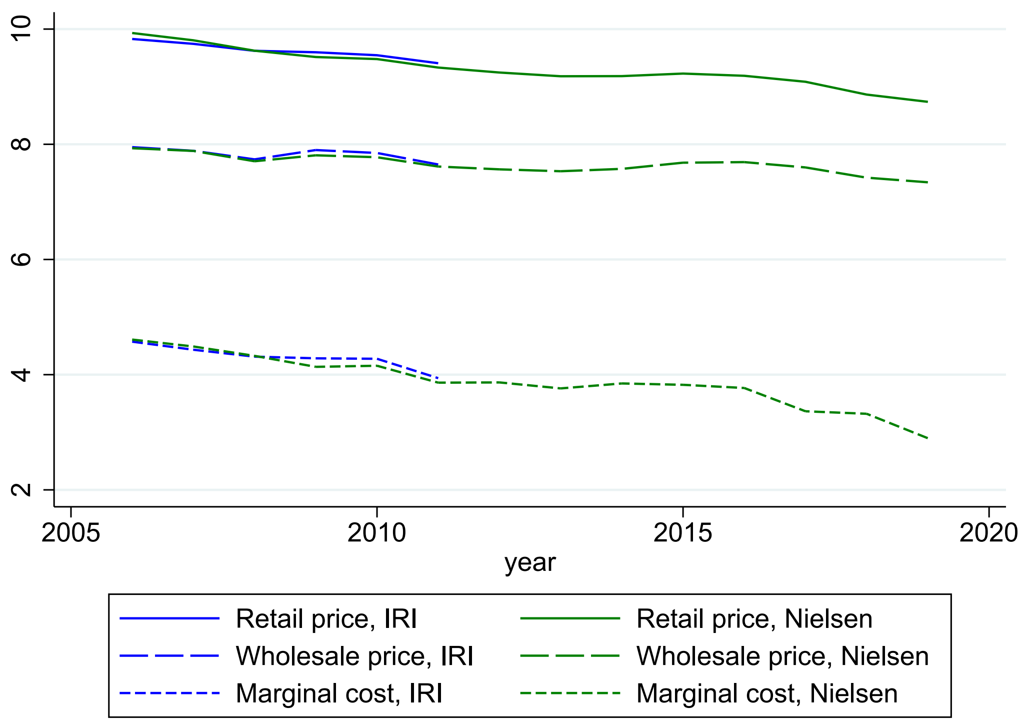

Notes Prices and marginal costs are volume-weighted averages per 144 ounces. Marginal costs are recovered using MW's demand system, assuming Nash-Bertrand competition, and correcting for downstream retail and wholesale costs as described in the text. Following MW, the merger period of June 2008 to May 2009 is omitted from these calculations. Dollar values deflated to 2010 levels using the Consumer Price Index.

Figure 4 is consistent with Ashenfelter et al. (2015), who find that the Miller-Coors merger led to cost efficiencies and increased markups that roughly offset each other. These time paths also mirror recent studies that find increasing markups driven by decreasing marginal cost, rather than increased prices (Brand, 2021; Döpper et al., 2021; Miller et al., 2023). In contrast, Grieco et al. (2024) find that the US automobile industry has evolved differently, with increasing prices and decreasing markups.

5 Combining production and demand analysis

When one can implement both the production- and demand-based approaches, markups are overidentified, and this overidentification can be used to relax or test modeling assumptions, both economic and econometric.

First, we introduce markups from the production approach as an additional moment in demand estimation, potentially replacing instruments used to identify demand. We impose a moment that is the difference between the production-based average brewer markup and the average brewer markup implied by the demand model (imposing a model of price determination). While Miller and Weinberg's cost shifter instruments are plausibly valid, such instruments are not readily available in all settings, and this exercise shows that production-based markups can be used to evaluate or supersede cost-shifters.

Second, we combine the two approaches to shed light on the form of competition in the retail market, illustrating how a joint (production and demand-conduct) approach allows one to verify otherwise ad hoc assumptions about market structure and firm conduct. We introduce a stylized conceptual framework that combines the two approaches to test the degree of market power enjoyed by retailers.

The first exercise combines the two approaches to evaluate or relax an *econometric* assumption; the second exercise, an *economic* assumption.

5.1 Estimating demand using production-based markups

In this section, we illustrate how production-based estimates can facilitate demand estimation. We revisit the estimation of Miller and Weinberg's demand system and estimate demand as if some of their instruments were unavailable (e.g., the cost shifters). We replace the missing instruments with a moment based on our production-based estimate of average markups. The resulting objective function penalizes demand parameters that imply mean markups that differ from our production-based estimates.[53]

Miller and Weinberg use four sets of instruments in their demand estimation: cost shifters based on distances between breweries and markets, instruments based on the number of products in the product set ("BLP instruments"), instruments based on demographics, and a dummy variable indicating Anheuser-Busch InBev and MillerCoors products in periods following the Miller-Coors merger. Appendix B includes a full description of these instruments.

We re-estimate MW's demand system dropping either their cost shifters or BLP instru-

[53] Conlon and Rao (2024) incorporate markup moments into demand estimation in a similar manner; they observe markups directly rather than estimating them with the production approach.

ments, and impose instead the moment

$$E\left[w_{jdt}\left(\frac{P_{jdt}}{P_{jdt}-c_{jdt}(\theta^D)}-\widehat{\mu}\right)\right]=0,$$

where $\widehat{\mu}$ is the mean markup estimate from our production analysis (2.076), w_{jdt} is a weight corresponding to revenue, and $c_{jdt}(\theta^D)$ is the marginal cost of product j in market d and period t (see equation 3), given demand parameters θ^D.[54] Intuitively, the production-based markup provides information about own-price elasticities, for markups are of course inversely proportional to own-price elasticities. Details of how we implement the estimation of MW's demand system with this moment can be found in Appendix B.

Table 4 shows that estimation of MW's demand system using the production-based markup moment results in demand systems that are broadly similar to MW's original estimates. Column 1 reports estimates of MW's main demand specification, which they call RCNL-1, using their original moments and code. Column 2 reports our estimates when dropping MW's cost shifters, and imposing instead the production-based markup moment. While the mean price coefficient in our estimates appears much larger in absolute value, the nesting parameter is lower, and these differences partially cancel, resulting in a mean own-price elasticity that is only slightly less elastic.[55] A more substantial difference is that our estimates feature more substitution to the outside option, a consequence of the smaller nesting parameter.

In column 3, we estimate the demand system with the markup moment again, but now dropping the BLP instruments (instruments involving the number of products). These estimates are quite similar to the version without cost shifters. Although cost shifters and BLP instruments play different roles in the identification of demand (Berry and Haile, 2014, 2016), the markup moment can replace either, at least in this case.

Columns 4 and 5 are included to confirm that the markup moment is playing an important role. Column 4 shows that when neither cost shifters nor the markup moment are included, the estimates are quite implausible. Most notably, the implied markups are extremely high.

In column 5, we perform the estimation without BLP instruments or the markup moment. In this case, the resulting demand system isn't obviously implausible, but it is quite different from MW's baseline estimates, much further away from the baseline than the estimates with the markup moment.

[54] We caution readers about using markups defined as p/mc in moments. Defined this way, the markup is not a well-behaved statistic, for it can asymptote to infinity and then flip signs as demand elasticitities change. We address this issue by top-coding markups at 15 as described in Section 4.2. Using marginal costs instead of markups in moments would also avoid moments that are discontinuous in parameters.

[55] A more comparable measure of price sensitivity across columns is $\alpha/(1-\rho)$, where α is the coefficient on price, and ρ is the nesting parameter. This varies little across columns, being equal to -0.513 for column 1 and -0.518 for column 2, for example.

Table 4: Demand Estimates With Markup Moments

	(1)	(2)	(3)	(4)	(5)
Price	-0.0832	-0.1943	-0.1945	-0.0535	-0.0456
	(0.0137)	(0.0451)	(0.0566)	(0.0005)	(0.0006)
Nesting Parameter	0.8378	0.6252	0.6245	0.8507	0.9264
	(0.0386)	(0.0737)	(0.1023)	(0.0057)	(0.0032)
Income × Price	0.0006	0.0019	0.0019	0.0006	0.0001
	(0.0002)	(0.0005)	(0.0006)	(0.0067)	(0.0071)
Income × Constant	0.0144	0.0062	0.0063	0.0171	0.0185
	(0.0049)	(0.0071)	(0.0081)	(0.0026)	(0.0029)
Income × Calories	0.0040	0.0069	0.0069	0.0025	0.0027
	(0.0015)	(0.0025)	(0.0027)	(0.5779)	(1.2817)
Mean Brewer Markup	2.2991	2.0764	2.0767	6.5734	1.7837
Median Own Price Elasticity	-4.7133	-4.2678	-4.2663	-2.8511	-6.4012
Median Market Price Elasticity	-0.5698	-1.1933	-1.1946	-0.3302	-0.3455
Median Outside Diversion	0.1240	0.2857	0.2862	0.1184	0.0559
Markup Moment	No	Yes	Yes	No	No
Distance Instruments	Yes	No	Yes	No	Yes
BLP Instruments	Yes	Yes	No	Yes	No
Demographic Instruments	Yes	Yes	Yes	Yes	Yes
Post-Merger Macrobrewery Dummy	Yes	Yes	Yes	Yes	Yes

Notes: Column 1 presents MW's demand system estimated with their original moments. Columns 2 and 3 drop some of their moments as indicated, and impose instead a moment based on a difference between the demand-based mean brewer markup and the production-based estimate of 2.076. Columns 4 and 5 drop some of MW's instruments without imposing the markup moment. Standard errors in parentheses take into account uncertainty in the production-based markup estimate. See Appendix B for details.

Including a moment based on production-based estimates is conceptually similar to including cost shifters and "supply moments" in demand estimation, following Berry et al. (1995) and much of the subsequent literature, but there is an important practical difference. Dropping the cost shifters in column 2, we no longer rely on the availability of an exogenous cost shifter. Alternatively, in column 3, we avoid the assumption that demand shocks are uncorrelated with the BLP instruments (in this context, the number of products). These are significant advantages. In general, strong and exogenous cost shifters are not easy to find—in our view, MW's instruments represent an unusually good case. Production function estimation may be feasible in many situations where valid cost shifters are not readily available, or when product sets (and BLP instruments) do not vary, and production-based moments can serve to check an instrument's validity. Therefore, we expect production analysis has the potential to serve as a complement to demand estimation quite broadly.

There are many other possibilities for how production analysis could be integrated into demand estimation. When production and demand data can be merged, a richer set of moments is feasible—e.g., one could stipulate that the difference between demand- and production-based markups should be uncorrelated with other observables, in addition to matching in levels. One could also use production analysis to assist with the identification of a demand model with richer consumer heterogeneity, more flexible substitution patterns, and more parameters.

In this section, the combined approaches allow us to avoid some exclusion restrictions—*econometric* assumptions. In the next, we will combine the approaches to test an *economic* assumption.

5.2 Identifying the degree of retail competition

In this section, we use the two approaches together to test the degree of competition in the retail market. To identify producer-level markups using the demand-based approach, one makes assumptions (implicitly or explicitly) regarding how retail markets operate. The most common approach in the literature, from Berry et al. (1995) to more recent demand-based studies of markups, including Brand (2021), Döpper et al. (2021), and Grieco et al. (2024), is to ignore the retail sector and suppose that retail prices are controlled by producers. Such an approach requires strong assumptions to justify. E.g., one could assume that the retail market is perfectly competitive or that producers engage in resale price maintenance. Even then, downstream costs can drive a wedge between wholesale and retail prices that will bias estimates of producers' markups if ignored.

We maintained the assumption of perfectly competitive retail markets above, although we did allow for costs of retailing to create a gap between wholesale and retail prices. Now, we relax the assumption of competitive retail with the aim of assessing whether retail market power should be accounted for when estimating producer-level markups using the demand approach. Although the (often implicit) assumption of price-taking retailers has become relatively standard when recovering markups using the demand approach, studies such as Hoch et al. (1994) and Chevalier (1995) provide evidence of retail market power that calls this practice into question.

We introduce a simple and stylized model of store choice, representing the degree of differentiation in the retail market through a single parameter. Our paper complements studies like Ellickson et al. (2020); while their study features a much richer model of store choice and is capable of assessing mergers between specific firms, ours parsimoniously characterizes the overall level of competitiveness in the retail sector with one parameter, naturally allows one to map to producer markups, and avoids the need to estimate the disutility of distance to store. A separate literature includes a store choice component in the context of studying the vertical contracts between retailers and producers (Berto Villas-Boas, 2007; Bonnet and Dubois, 2010), but these papers do not provide direct evidence on the degree of competition in the retail sector.[56]

At the core of our joint production-demand approach lies the distinction between *store-level demand* and a *product demand* system. A store-level demand system describes what happens at an individual store when it changes its prices, holding prices elsewhere fixed. A product demand system describes the behavior of a consumer (or consumers) when they are faced with a menu of products and a vector of associated prices, abstracting away from store choice. If one ignores the presence of retailers, and supposes that producers sell directly to consumers, like Berry et al. (1995) and much of the following literature, then the product demand system is what producers consider when setting prices.

These are two distinct demand systems, only equal to each other when retailers have monopoly power over their shoppers. It is precisely the degree of retail competition, and the form of competition in the relevant market, that we treat as unknown and to be recovered from the data.

Given brewer markups from the production approach, we can back out the degree of retail competition that produces similar demand-based brewer markups. In principle, this approach can be applied by product-market-time, however, given the restrictions we face using confidential census data, we adopt an aggregate version of the approach (across markets and products). Finally, it should come as no surprise that in order to recover information about conduct in the retail market, more structure is required. Throughout we illustrate this joint approach by relying on a symmetric equilibrium among retailers.

Before we develop the analytical framework, it is important to underscore that the

[56] To be more precise, Berto Villas-Boas (2007) and Bonnet and Dubois (2010), two similar products in one store are differentiated in the same way as two similar products at different stores—the parametric model imposes similar price responsiveness within-store and across stores.

production approach does not rely on any assumptions regarding either brewer or retail competition or conduct. Conditional on consistently estimating the production function and on having variable input(s) that are chosen to minimize costs, the production approach delivers consistent estimates of plant-level markups regardless of how downstream competition plays out.

While we combine the demand and production approaches here to shed light on the nature of beer retail competition, we believe that the approaches could be combined in many fruitful ways. Section 5.2.4 discusses other possibilities.

5.2.1 Retail competition and producer markups

We start by developing the main intuition behind our approach using a simplified framework. This will serve to highlight that, in general, when the vertical structure of industries is taken into account, there will be a wedge between a producer's residual demand elasticity (which features in the producer's first-order condition to set prices) and the elasticities from the product demand system. Starting from the premise that we observe brewer markups and the product demand system, we present a model of retail competition that allows us to isolate retail competitiveness.

Consider a retail market consisting of symmetric retailers with identical costs and product offerings. To simplify the exposition, we assume that there is only one brewer, which produces a single product. The intuition we develop in this section generalize to multiple brewers with multiple products, and our empirical approach below allows for such.

As before, P^r denotes the price of beer charged to consumers by retailer r; P denotes the wholesale price charged by the brewer to retailers (we assume brewers do not price discriminate within a retail market). Let q^r be the quantity sold by a particular retailer, and let $Q = \sum_r q^r$ be the total market-level quantity.[57]

We assume that retailers are in a symmetric equilibrium of a price setting game and we denote the equilibrium retail price by $P^R(P)$. That is, with two retailers and wholesale price P, *both* retailers charge price $P^R(P)$ in equilibrium. We use a capital R rather than lowercase in $P^R(\cdot)$ to emphasize that the function describes the symmetric equilibrium price, rather than an individual retailer's price.

A key observation is that the brewer, in general, does not face the same residual demand elasticity as does a retailer. This is best illustrated by writing out the brewer's elasticity:

$$\begin{aligned}\eta &= \frac{\partial Q}{\partial P^R}\frac{dP^R(P)}{dP}\frac{P}{Q} \\ &= \frac{\partial Q}{\partial P^R}\frac{dP^R(P)}{dP}\frac{P}{P^R}\frac{P^R}{Q} = \eta^R\frac{dP^R(P)}{dP}\frac{P}{P^R},\end{aligned} \tag{17}$$

[57] We suppress the market-period subscript t.

where η^R is the elasticity of *total* retail demand with respect to a change in the equilibrium retail price (i.e., a change in the price charged by all retailers), and $\frac{dP^R(P)}{dP}$ represents wholesale-retail pass-through. Demand faced by the brewer is less elastic than retail demand as long as $\frac{dP^R(P)}{dP} \leq 1$, given that $\frac{P}{P^R} \leq 1$. Even with full pass-through $\left(\frac{dP^R(P)}{dP} = 1\right)$, the gap between retail and wholesale prices will make upstream demand less elastic than retail demand.

In many contexts, including the Miller-Weinberg demand system presented above, demand estimation yields a *product demand system*—a demand system that describes choice within a product set without any store choice. The retail elasticity η^R can be understood as a feature of the product-level demand system, for as we change symmetric retailers' prices together, store choice is unaffected.

Thus, η^R is effectively known after estimating the product demand system. Equation (17) then shows how the elasticity faced by the brewer differs from this product-demand elasticity. An important part of the difference is the (inverse) retailer's markup $\frac{P}{P^R}$. This margin will be related to the retailer's first-order condition for pricing:

$$P^r = (P + c^r)\left[1 + \frac{1}{\eta^r}\right]^{-1}, \tag{18}$$

where c^r is the marginal cost of retailing, and η^r is an individual retailer's demand elasticity with respect to a unilateral price change, which is potentially much more elastic than η^R. The distinction between η^r and η^R is crucial for this exercise. The former describes what happens when one retailer changes their prices; the latter describes what happens when all retailers implement a similar price change.

When the retail environment is very competitive, the elasticity η^r will be very large in absolute value, there will be little markup, and the gap between the brewer's price P and the retail price P^r will correspond to the marginal cost of retailing. When retailers have considerable market power, η^r will be more modest (in the limit of retail monopoly, $\eta^r = \eta^R$), and retail margins will be more substantial.

Putting the above points together, retail margins create a wedge between the demand elasticity faced by brewers and the (estimated) elasticity of the product demand system, and this wedge is larger when retailers have more market power. Therefore, brewer markups recovered by the demand approach are a function of the competitiveness of the retail environment.

In the following section, we introduce a specific way of parameterizing the degree of competition among retailers, and we implement this in the context of our setting and data. Taking the Miller-Weinberg demand system as our product demand system, we can then derive retail margins and retail-wholesale pass-through as a function of the retail competition parameter.

5.2.2 Full model of retail competition

We now introduce the empirical framework we take to the data. We consider two symmetric but differentiated retailers. A parameter $\gamma \in [0,1]$ captures the degree of differentiation between the two. While the assumption of two symmetric retailers is stylized, it is flexible in an important sense, as this setup will allow us to span the space between retail monopoly and perfectly competitive retail. Ultimately, this structure allows us to derive retailer elasticities (e.g., η^r above) from the product demand system and γ, a parameter capturing the degree of differentiation among retailers.

Retail competition We assume that the vector of purchased quantities at a store is equal to

$$\boldsymbol{q}_t^r\left(\boldsymbol{P}^r, \boldsymbol{P}^{r'}\right) = N_t \int \widetilde{s}\left(\boldsymbol{P}^r, \boldsymbol{P}^{r'}; \theta\right) \boldsymbol{s}\left(\boldsymbol{P}^r; \theta\right) dF_t\left(\theta\right),$$

where $\boldsymbol{s}\left(\boldsymbol{P}^r; \theta\right)$ is the vector of market shares from Miller and Weinberg's demand system for consumer type θ, and $\widetilde{s}$ describes the share of shoppers in the area (note the t subscript here can pick up both the time period and the geographic market that the store operates within). $\boldsymbol{P}^{r'}$ represents the prices at competing stores and we assume that stores are in a symmetric equilibrium such that $\boldsymbol{P}^r = \boldsymbol{P}^{r'}$ in equilibrium, but we must distinguish between the two so that we can consider unilateral price changes by one store.

We assume that store choice is driven by a logit model of store-level inclusive values. That is,

$$\widetilde{s}\left(\boldsymbol{P}^r, \boldsymbol{P}^{r'}; \theta\right) = \frac{\exp\left(\frac{\gamma}{1-\gamma} V\left(\boldsymbol{P}^r; \theta\right)\right)}{\exp\left(\frac{\gamma}{1-\gamma} V\left(\boldsymbol{P}^r; \theta\right)\right) + \exp\left(\frac{\gamma}{1-\gamma} V\left(\boldsymbol{P}^{r'}; \theta\right)\right)}, \tag{19}$$

where $V\left(\boldsymbol{P}^r; \theta\right)$ is the inclusive value of a store's choice set with price vector $\boldsymbol{P}^r$ for a consumer of type θ, and γ is the parameter dictating the degree of retail competition; $\gamma = 0$ corresponds to retail monopoly—the share of customers coming to a store does not depend on prices—while as $\gamma \to 1$, we approach a perfectly competitive retail environment.

This model of store choice fits naturally with MW's estimated demand system. Their focus is on the demand system from the perspective of brewers, and their demand estimation aims explicitly at recovering a *product-level* demand system—indeed, their estimation relies on data that is aggregated across stores. Our setup allows us to take their estimates of the product-level demand system for granted, and we introduce the γ parameter to represent the degree of differentiation in the retail environment.

Retailers solve a standard profit maximization problem (suppressing the t subscript now for the sake of exposition):

$$\max_{\boldsymbol{P}^r} \left(\boldsymbol{P}^r - \boldsymbol{P} - c^r\right) \boldsymbol{q}^r\left(\boldsymbol{P}^r, \boldsymbol{P}^{r'}\right).$$

The retailer's first-order condition is

$$\frac{\partial \boldsymbol{q}^r\left(\boldsymbol{P}^r, \boldsymbol{P}^{r'}\right)}{\partial \boldsymbol{P}^r}\left(\boldsymbol{P}^r - \boldsymbol{P} - c^r\right) + \boldsymbol{q}^r\left(\boldsymbol{P}^r, \boldsymbol{P}^{r'}\right) = 0. \tag{20}$$

Importantly, we differentiate here with respect to the first argument of $\boldsymbol{q}^r\left(\boldsymbol{P}^r, \boldsymbol{P}^{r'}\right)$, but not the second – this represents one store making price changes, holding the prices of other stores fixed.

Ultimately, the γ parameter serves to create a wedge between retail prices and retail costs (the wholesale price and any retail costs). At $\gamma = 1$, we are in the extreme case where retail prices equal retail marginal costs, which in turn are equal to the cost of retailing plus wholesale prices charged by brewer-distributors. As we move away from $\gamma = 1$, we assign an increasing degree of market power to retailers, and this increases their markups (and the wedge between retail and wholesale prices).

We use equation (20) to recover the vector of brewer-distributor prices $\boldsymbol{P}$ conditional on γ. The derivative in equation (20) is a function of the estimated product demand system and retail competitiveness parameter, γ. The other terms are either observed in the data (retail prices and quantities) or imputed (the cost of retailing, which was explained in Section 3.1.2).

We can let $\boldsymbol{P}^R(\boldsymbol{P})$ denote the symmetric equilibrium retail price vector as a function of brewer prices $\boldsymbol{P}$. Such prices represent a solution to (20) with $\boldsymbol{P}^{r'} = \boldsymbol{P}^r$. We need not actually solve for such an equilibrium, so we merely assume the existence of such a price vector; uniqueness is not necessary. Ultimately, what we need is that $\boldsymbol{P}^R(\boldsymbol{P})$ is a differentiable function.

Brewer competition Before addressing the brewer's problem, we consider how brewer prices pass through to retail prices. We do this by applying the implicit function theorem to equation (20), evaluated at $\boldsymbol{P}^{r'} = \boldsymbol{P}^r$, with both equal to the symmetric equilibrium price vector $\boldsymbol{P}^R(\boldsymbol{P})$. It is important to emphasize that we impose a symmetric equilibrium ($\boldsymbol{P}^{r'} = \boldsymbol{P}^r$) *before* applying the implicit function theorem to (20), but after taking the price derivative expressed in (20). When the brewer changes wholesale prices, all retailers will adjust prices—that is, the retail equilibrium will adjust—and we are interested in the equilibrium response to a change in wholesale prices. This results in the following expression for retail-wholesale pass-through:

$$\begin{aligned}\frac{d\boldsymbol{P}^R(\boldsymbol{P})}{d\boldsymbol{P}} = \; & \Bigg[\left(\sum_{j=1}^{J}\frac{\partial}{\partial \boldsymbol{P}^r}\left(\left.\frac{\partial q_j^r\left(\boldsymbol{P}^r,\boldsymbol{P}^{r'}\right)}{\partial \boldsymbol{P}^r}\right|_{\boldsymbol{P}^{r'}=\boldsymbol{P}^r}\right)\right)\left(\boldsymbol{P}^r - \boldsymbol{P} - c^r\right) \\ & + \left(\frac{\partial \boldsymbol{q}^r\left(\boldsymbol{P}^r,\boldsymbol{P}^{r'}\right)}{\partial \boldsymbol{P}^r}\right)' + \frac{d\boldsymbol{q}^r\left(\boldsymbol{P}^r,\boldsymbol{P}^r\right)}{d\boldsymbol{P}^r}\Bigg]^{-1}\left(\frac{\partial \boldsymbol{q}^r\left(\boldsymbol{P}^r,\boldsymbol{P}^{r'}\right)}{\partial \boldsymbol{P}^r}\right)',\end{aligned}$$

which we evaluate with $\boldsymbol{P}^r$ and $\boldsymbol{P}^{r'}$ equal to the observed retail price vector. In understanding this equation, it is important to distinguish between the derivatives

$$\frac{d\boldsymbol{q}^r\left(\boldsymbol{P}^r,\boldsymbol{P}^r\right)}{d\boldsymbol{P}^r}$$

and

$$\frac{\partial\boldsymbol{q}^r\left(\boldsymbol{P}^r,\boldsymbol{P}^{r'}\right)}{\partial\boldsymbol{P}^r}.$$

The former involves changing prices at a store and its competitor in the same way, thus introducing no changes in store-level market shares – it corresponds to the derivative of the product-level demand system.[58] The latter is a partial derivative, changing a store's own prices while holding its competitor's prices fixed. The second derivatives of $\boldsymbol{q}^r$ involve derivatives of both product shares $\boldsymbol{s}$ and derivatives of store shares $\widetilde{s}$. The model of store shares, equation (19), allows us to compute the latter.

The pass-through matrix $\frac{d\boldsymbol{P}^R(\boldsymbol{P})}{d\boldsymbol{P}}$ converges to an identity matrix as $\gamma\rightarrow 1$ (as retail becomes perfectly competitive). This can be seen from the fact that the Jacobian term $\frac{\partial\boldsymbol{q}^r\left(\boldsymbol{P}^r,\boldsymbol{P}^{r'}\right)}{\partial\boldsymbol{P}^r}$ grows without bound as the market becomes more competitive. The other terms remain bounded, and so the above expression converges to the Jacobian term times its inverse.

We now let $\boldsymbol{q}\left(\boldsymbol{P}^r\right)\equiv\boldsymbol{q}^r\left(\boldsymbol{P}^r,\boldsymbol{P}^r\right)$ denote the vector of quantities sold in a symmetric retail equilibrium with retail price vector $\boldsymbol{P}^r$.[59] Accordingly, $\boldsymbol{q}\left(\boldsymbol{P}^R\left(\boldsymbol{P}\right)\right)$ gives equilibrium quantities as a function of wholesale prices. We can then write the brewer's pricing problem as follows:

$$\max_{\boldsymbol{P}_b}\left(\boldsymbol{P}_b-\boldsymbol{c}_b-c_b^w-\tau\right)'\boldsymbol{q}_b\left(\boldsymbol{P}^R\left(\boldsymbol{P}_b,\boldsymbol{P}_{-b}\right)\right),$$

where $\boldsymbol{P}_b$ is the vector of brewer prices for the products brewer b owns, $\boldsymbol{P}_{-b}$ are the brewer prices for other products, and $\boldsymbol{q}_b$ includes the elements of $\boldsymbol{q}$ corresponding to b's products.

We can write a brewer's first-order conditions as.

$$\left(O_b\frac{d\boldsymbol{q}\left(\boldsymbol{P}^r\right)}{d\boldsymbol{P}^r}\frac{d\boldsymbol{P}^R\left(\boldsymbol{P}\right)}{d\boldsymbol{P}}O_b'\right)'\left(\boldsymbol{P}_b-\boldsymbol{c}_b-c_b^w-\tau\right)+\boldsymbol{q}_b\left(\boldsymbol{P}\right)=\boldsymbol{0},$$

where O_b is a $J_b\times J$ ownership matrix with one element of each row equal to 1 to indicate products that b owns, and other elements equal to zero.

[58]More precisely, $\frac{d\boldsymbol{q}^r\left(\boldsymbol{P}^r,\boldsymbol{P}^{r'}\right)}{d\boldsymbol{P}^r}$ is equal to the derivative of the product demand system $\frac{d\boldsymbol{s}(\boldsymbol{P}^r)}{d\boldsymbol{P}^r}$ times the store's share (.5) and market population.

[59]Note that this $\boldsymbol{q}$ function is just the product demand system rescaled:

$$\boldsymbol{q}_t\left(\boldsymbol{P}^r\right)=N_t\int .5\boldsymbol{s}\left(\boldsymbol{P}^r;\theta\right)dF_t\left(\theta\right).$$

We can then recover marginal cost as follows:

$$\boldsymbol{c}_b = \boldsymbol{P}_b - c_b^w - \tau + \left[\left(O_b \frac{d\boldsymbol{q}\left(\boldsymbol{P}^r\right)}{d\boldsymbol{P}^r}\frac{d\boldsymbol{P}^R\left(\boldsymbol{P}\right)}{d\boldsymbol{P}} O_b'\right)'\right]^{-1} \boldsymbol{q}_b\left(\boldsymbol{P}\right). \tag{21}$$

Before applying equation 21, we recover brewer-distributor prices $\boldsymbol{P}$ (of which $\boldsymbol{P}_b$ is a subvector) using equation (20). Then, applying equation (21) for each brewer, we can recover the full vector of brewers' marginal costs. When applying this to data, $\boldsymbol{q}\left(\boldsymbol{P}\right)$ corresponds to observed quantities. Derivatives should also be evaluated at the observed prices. Brewer markups may then be calculated as before (see equation (5)).

We apply this for each market in 2007 and each domestic producer present in MW's demand system.

5.2.3 Results

Figure 5 presents revenue-weighted mean brewer markups implied by the above approach. It shows that for high values of the retail competitiveness parameter γ, demand-based markups are close to the production based markups. However, as we move away from a competitive retail environment, the markups from the two approaches quickly diverge, and for values of $\gamma \leq .97$, we have no overlap in the 95% confidence intervals from the two approaches.

If we take for granted the MW estimates of the product-level demand system as well as our production-based markup estimates, these results imply meaningful bounds on the degree of retail competition. As illustrated by Figure 6, retailer markups are about 1.3 at $\gamma = .97$. This retail markup refers to the ratio of retail price to retailer's full marginal cost, including costs of retailing and the brewer-distributor's wholesale price (including the brewer's price received, distribution costs, and excise taxes).

Figure 6 also serves to illustrate *why* demand-based brewer markups increase as we move away from a perfectly competitive retail environment. Recalling equation (17), differences between the demand elasticities faced by brewers and the demand elasticities of the product demand system come from two sources: incomplete pass-through and retail markups. Figure 6 shows that retail markups increase substantially as the retail environment becomes less competitive. In contrast, incomplete pass-through has a relatively small impact on the demand elasticities faced by brewers. The minimum and maximum rates of pass-through remain close to unity even as retailer markups become quite large. This breakdown indicates that our results are compatible with complete pass-through by retailers with a modest retail markup (of as much as $\approx 30\%$ of the retailers cost), even if such behavior results from a very different model of retail price determination than ours.

This rejection of high levels of retail market power is consistent with the findings in Miller and Weinberg's 2017 Appendix E. What distinguishes our results is the nature of the

Figure 5: Brewer markups and retail market structure

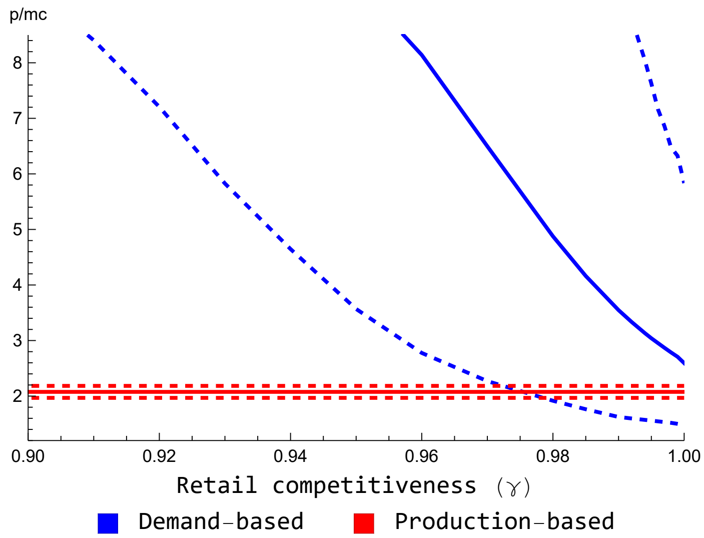

Notes: The blue lines represent implied markups for values of the retail competitiveness parameter (γ), and the dashed lines indicate the 95% confidence interval (obtained by simulating the asymptotic distribution of demand estimates). The solid red line represents the estimated (average) markup using the production approach (using the Leontief-Cobb-Douglas specification), while the dashed red lines represent its 95% confidence interval.

evidence. While past studies of retail conduct have been contained to analysis of demand, we rely on information about the level of brewer markups based on production analysis.

A potential concern is that our model of downstream costs included a measure of costs of retailing that were based on observed differences between retail and wholesale prices in the Dominick's Finer Foods data. If retailers have market power, the gap between retail and wholesale prices will correspond to more than just costs. However, this imputation has a relatively small impact on our results. Figure C3 includes a specification where costs of retailing are assumed to be zero; qualitatively, our findings are unchanged.

We cannot reject the null hypothesis of passive (perfectly competitive) retailers, and our results cast doubt on the possibility that retailers have considerable market power when selling beer. That said, they could have *some nonzero market power*. In light of studies like Hoch et al. (1994) and Chevalier (1995), we should distinguish between the question "are retail markets literally perfectly competitive?" and "can we take seriously demand-based estimates of producers markups that do not account for retail market power?" While the answer to the former is undoubtedly negative, our results are consistent with an affirmative answer to the latter, at least in the context of beer.

Nonetheless, competitive retail markets are compatible with a large gap between retail

Figure 6: Retail margins and pass-through

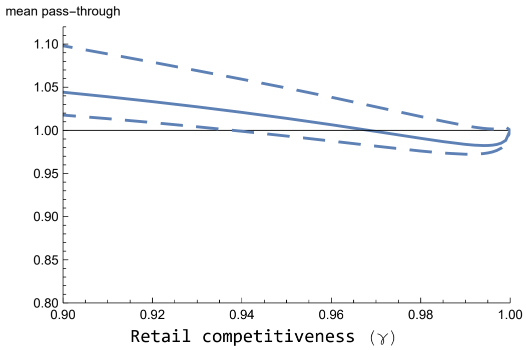

(a) Mean wholesale-retail pass-through

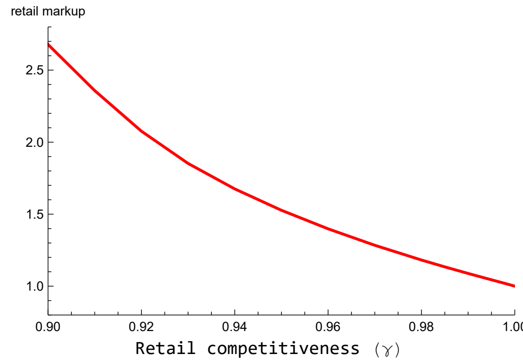

(b) Mean retail makup

Notes: Panel (a) displays the revenue-weighted mean rate of pass-through of wholesale to retail prices, that is, the average diagonal element of $d\boldsymbol{P}^r/d\boldsymbol{P}$. Dashed lines represent the maximum and minimum values of pass-through. We take the mean across markets after taking the max or min over diagonal elements of $d\boldsymbol{P}^r/d\boldsymbol{P}$ within each market. Panel (b) displays the share-weighted mean markup for retailers. For these markups, the numerator is the retail price; the denominator is the brewer's price plus distribution costs, excise taxes, and retailing costs.

and producer prices, for that gap can result from downstream costs even without downstream market power. Therefore, downstream costs (and the resulting retail-wholesale price gap) are an important factor to take into account when recovering markups of producers using retail demand data. As Table 1 indicates, omitting the costs of retail and/or distribution makes a substantial difference in the producer markups we estimate, a potentially important consideration in the context of demand-based studies of changes in producer's markups over time ((Brand, 2021; Grieco et al., 2024; Döpper et al., 2021)) that do not allow for changes in downstream costs over time.

5.2.4 Other conduct tests

In this section, we have used information about upstream markups to study downstream market power, but it would be natural to use producer markups to study producer conduct. For example, our production-based estimates allow us to add to MW's conclusions about brewer conduct. MW argue that the US brewing industry was probably not in Nash-Bertrand equilibrium both before and after the Miller-Coors merger. One interpretation of their evidence is that Nash-Bertrand equilibrium prevailed before the merger, and then a form of tacit collusion prevailed after. However, MW cannot actually verify that the industry was in Nash-Bertrand equilibrium before the merger; "The nature of pre-merger competition is not determined" (p. 1784). In contrast, we can provide direct evidence on the plausibility of Nash-Bertrand conduct before the merger—demand-based markups implied by Nash-Bertrand competition are indeed consistent with our production based markup estimates in 2007 (before the merger), supporting the hypothesis that pre-merger conduct was Nash-Bertrand.

Putting aside the distinction between upstream and downstream conduct, our study of the mode of competition follows the spirit of Bresnahan (1987), where conduct is tested by contrasting the different models of competition against each other, given the estimated demand system. While Bresnahan (1982) and Berry and Haile (2014) show that testing conduct is possible within the demand approach, their results rely on particular variation in the data (e.g., demand shifters that can be excluded from marginal cost functions).[60] In contrast, our approach, which appeals to production-based estimates, allows us to test conduct based on the levels of markups recovered from production data. Intuitively, the production approach offers measures of marginal cost, and directly observing marginal cost makes it easy to evaluate models of conduct. Thus, we expect production analysis combined with demand estimation could provide a fruitful way to study producer conduct.

6 Conclusion

While inferring markups from demand data is common practice, the method relies on assumptions on consumer choice and how firms compete in a market. Alternatively, markups can be inferred from production data, relying on the assumptions about the timing of inputs and evolution of productivity.

In this paper, we compare markup estimates from both strategies for the case of the US brewing industry, following standard practice in each case. We find the two approaches agree on the level of average markups in the US Brewing industry, both in terms of the overall level and regarding a recent upward trend. In particular, we find that markups rise, on average, from below 1.5 in 1972 to over 2 in recent years.

We then show how combining the approaches can serve to test or relax both economic and econometric assumptions. First, we show that information on markups from production analysis can serve as a substitute for cost-shifting instruments in demand estimation. Thus, combining the approaches allows us to relax an exclusion restriction used in demand estimation. Next, we combine the approaches to test the level of retail competition. Our estimates indicate that retail competition is approximated well by perfect competition, or at least with passive retailers that fully pass through wholesale price changes to retail prices. This is at first reassuring news for demand-based studies of producer-level markups and marginal cost that treat the retail sector as non-existent, perfectly competitive, or otherwise imposing perfect wholesale-retail pass-through. Thus, combining the approaches allows

[60]See Duarte et al. (2021) and Backus et al. (2021) for implementations.

us to test an economic assumption used to recover markups from demand estimates.

Structural approaches to studying market power have been criticized by Angrist and Pischke (2010), in essence, for relying on a host of difficult-to-test assumptions. The results from this paper show that two very different approaches give compatible results on the degree of market power in the US beer industry, and they suggest that combining the demand and production approaches may assist researchers in selecting assumptions and evaluating structural models.

References

Ackerberg, D., L. Benkard, S. Berry, and A. Pakes (2007). Econometric tools for analyzing market outcomes. *The Handbook of Econometrics*, Chapter 63.

Ackerberg, D., R. Caves, and G. Frazier (2015, November). Identification properties of recent production function estimators. *Econometrica 83*(6), 2411–2451.

Ackerberg, D. and J. Hahn (2015, December). Some non-parametric identification results using timing and information set assumptions.

Ackerberg, D. A., K. Caves, and G. Frazer (2015). Identification properties of recent production function estimators. *Econometrica 83*(6), 2411–2451.

Ackerberg, D. A. and J. De Loecker (2024). Production function identification under imperfect competition. Technical report, Mimeo, University of Leuven.

Alviarez, V., K. Head, and T. Mayer (2020). Global giants and local stars: How changes in brand ownership affect competition.

Angrist, J. and J.-S. Pischke (2010). The credibility revolution in empirical economics: How better research design is taking the con out of econometrics. *Journal of Economic Perspective 24*(2), 3–30.

Arkolakis, C., A. Costinot, D. Donaldson, and A. Rodriguez-Clare (2015). The elusive pro-competitive effects of trade. *MIT mimeo*.

Ashenfelter, O. C., D. S. Hosken, and M. C. Weinberg (2015). Efficiencies brewed: pricing and consolidation in the us beer industry. *The RAND Journal of Economics 46*(2), 328–361.

Asker, J. (2016). Diagnosing foreclosure due to exclusive dealing. *The Journal of Industrial Economics 64*(3), 375–410.

Asker, J. and V. Nocke (2021). Collusion, mergers, and related antitrust issues. In *Handbook of Industrial Organization*, Volume 5, pp. 177–279. Elsevier.

Backus, M., C. Conlon, and M. Sinkinson (2021). Common ownership and competition in the ready-to-eat cereal industry. Technical report, National Bureau of Economic Research.

Berry, S., M. Gaynor, and F. Scott Morton (2019). Do increasing markups matter? lessons from empirical industrial organization. *Journal of Economic Perspectives 33*(3), 44–68.

Berry, S. and P. Haile (2016). Identification in differentiated products markets. *Annual review of Economics 8*(1), 27–52.

Berry, S., J. Levinsohn, and A. Pakes (1995, July). Automobile prices in market equilibrium. *Econometrica 63*(4), 841–90.

Berry, S., J. Levinsohn, and A. Pakes (2004). Differentiated products demand systems from a combination of micro and macro data: The new car market. *The Journal of Political Economy 112*(1), 68–105.

Berry, S. T. and P. A. Haile (2014). Identification in differentiated products markets using market level data. *Econometrica 82*(5), 1749–1797.

Berry, S. T. and P. A. Haile (2021). Foundations of demand estimation. In *Handbook of Industrial Organization*, Volume 4, pp. 1–62. Elsevier.

Berto Villas-Boas, S. (2007). Vertical relationships between manufacturers and retailers: Inference with limited data. *The Review of Economic Studies 74*(2), 625–652.

Bet, G. (2021). Market power in the us airline industry. *Available at SSRN 3913695*.

Bonnet, C. and P. Dubois (2010). Inference on vertical contracts between manufacturers and retailers allowing for nonlinear pricing and resale price maintenance. *The RAND Journal of Economics 41*(1), 139–164.

Brand, J. (2021). Differences in differentiation: Rising variety and markups in retail food stores. *Available at SSRN 3712513*.

Bresnahan, T. F. (1982). The oligopoly solution concept is identified. *Economics Letters 10*(1), 87–92.

Bresnahan, T. F. (1987). Competition and collusion in the american automobile industry: The 1955 price war. *The Journal of Industrial Economics*, 457–482.

Bronnenberg, B. J., M. W. Kruger, and C. F. Mela (2008). Database paper—the iri marketing data set. *Marketing science 27*(4), 745–748.

Chevalier, J. A. (1995). Capital structure and product-market competition: Empirical evidence from the supermarket industry. *The American Economic Review*, 415–435.

Collard-Wexler, A. and J. De Loecker (2016). Production function estimation with measurement error in inputs. Technical report, National Bureau of Economic Research.

Conlon, C. and J. Gortmaker (2020). Best practices for differentiated products demand estimation with pyblp. *The RAND Journal of Economics 51*(4), 1108–1161.

Conlon, C. and N. L. Rao (2024). The cost of curbing externalities with market power: Alcohol regulations and tax alternatives. Technical report, National Bureau of Economic Research.

Consumer Reports (1996, June). Can you judge a beer by its label? *Consumer Reports*, 10–12.

De Loecker, J. (2011). Product differentiation, multi-product firms and estimating the impact of trade liberalization on productivity. *Econometrica 79*(5), 1407—1451.

De Loecker, J. (2011). Recovering markups from production data. *International Journal of Industrial Organization 29*, 350–355.

De Loecker, J., J. Eeckhout, and G. Unger (2020). The rise of market power and the macroeconomic implications. *The Quarterly Journal of Economics 135*(2), 561–644.

De Loecker, J., P. K. Goldberg, A. K. Khandelwal, and N. Pavcnik (2016). Prices, markups, and trade reform. *Econometrica 84*(2), 445–510.

De Loecker, J. and P. T. Scott. Estimating market power: Evidence from the US brewing industry. NBER Working Paper.

De Loecker, J. and C. Syverson (2021). An industrial organization perspective on productivity. In *Handbook of Industrial Organization*, Volume 4, pp. 141–223. Elsevier.

De Loecker, J. and F. Warzynski (2012). Markups and firm-level export status. *American Economic Review 102*(6), 2437–71.

Dobbelaere, S. and J. Mairesse (2013). Panel data estimates of the production function and product and labor market imperfections. *Journal of Applied Econometrics 28*(1).

Döpper, H., A. MacKay, N. Miller, and J. Stiebale (2021). Rising markups and the role of consumer preferences. *Available at SSRN 3939126*.

Doraszelski, U. and J. Jaumandreu (2021). Reexamining the de loecker & warzynski (2012) method for estimating markups.

Duarte, M., L. Magnolfi, M. Sølvsten, and C. Sullivan (2021). Testing firm conduct. Technical report, working paper, University of Wisconsin-Madison.

Ellickson, P. B., P. L. Grieco, and O. Khvastunov (2020). Measuring competition in spatial retail. *The RAND Journal of Economics 51*(1), 189–232.

Goldberg, P. K. and R. Hellerstein (2013). A structural approach to identifying the sources of local-currency price stability. *Review of Economic Studies 80*(1), 175–210.

Grieco, P. L., C. Murry, J. Pinkse, and S. Sagl (2022). Conformant and efficient estimation of discrete choice demand models.

Grieco, P. L., C. Murry, and A. Yurukoglu (2024). The evolution of market power in the us automobile industry. *The Quarterly Journal of Economics 139*(2), 1201–1253.

Hall, R. E. (1988). The relation between price and marginal cost in us industry. *Journal of Political Economy 96*(5), 921–947.

Hellerstein, R. (2008). Who bears the cost of a change in the exchange rate? pass-through accounting for the case of beer. *Journal of International Economics 76*(1), 14 – 32.

Hidalgo, J. (2023). This craft's for you! entry and market (ing) competition in the us beer industry.

Hoch, S. J., X. Dreze, and M. E. Purk (1994). Edlp, hi-lo, and margin arithmetic. *Journal of marketing 58*(4), 16–27.

Hsieh, C.-T. and P. J. Klenow (2009). Misallocation and manufacturing TFP in china and india. *The Quarterly Journal of Economics 124*(4), 1403–1448.

Kirov, I. and J. Traina (2023). Labor market power and technological change in us manufacturing. Technical report, Working Paper.

Klette, T. J. and Z. Griliches (1996). The inconsistency of common scale estimators when output prices are unobserved and endogenous. *Journal of Applied Econometrics 11*(4), 343–361.

Lee, R. S., M. D. Whinston, and A. Yurukoglu (2021). Structural empirical analysis of contracting in vertical markets. In *Handbook of Industrial Organization*, Volume 4, pp. 673–742. Elsevier.

Levinsohn, J. and A. Petrin (2003). Estimating production functions using inputs to control for unobservables. *Review of Economic Studies 70*(2), 317–341.

Marschak, J. and J. Andrews, William H. (1944). Random simultaneous equations and the theory of production. *Econometrica 12*(3/4), 143–205.

Miller, N., M. Osborne, G. Sheu, and G. Sileo (2023). Technology and market power: The united states cement industry, 1974-2019. *Georgetown McDonough School of Business Research Paper* (4041168).

Miller, N. H. (2024). Industrial organization and the rise of market power. Technical report, National Bureau of Economic Research.

Miller, N. H. and M. C. Weinberg (2017). Understanding the price effects of the millercoors joint venture. *Econometrica 85*(6), 1763–1791.

Morlacco, M. (2020). Market power in input markets: Theory and evidence from french manufacturing. mimeo USC.

Nelson, J. (2005). Beer advertising and marketing update: Structure, conduct, and social costs. *Review of Industrial Organization 26*, 269–306.

Nevo, A. (2001). Measuring market power in the ready-to-eat cereal industry. *69*(2), 307–342.

Olley, G. S. and A. Pakes (1996). The dynamics of productivity in the telecommunications industry. *Econometrica 64*(6), 1263–1297.

Orr, S., P. Morrow, S. Rachapalli, and J. Cairncross (2024). Identifying firm vs. product markups using production data: Micro estimates and aggregate implications. *Product Markups Using Production Data: Micro Estimates and Aggregate Implications (July 23, 2024)*.

Peeters, M. (2014). Heterogeneous mark-ups, growth and endogenous misallocation. *mimeo Yale University*.

Raval, D. (2023). Testing the production approach to markup estimation. *Review of Economic Studies 90*(5), 2592–2611.

Rojas, C. (2008). Price competition in us brewing. *The Journal of Industrial Economics 56*(1), 1–31.

Rosse, J. N. (1970). Estimating cost function parameters without using cost data: Illustrated methodology. *Econometrica: Journal of the Econometric Society*, 256–275.

Rubens, M. (2023). Market structure, oligopsony power, and productivity. *American Economic Review 113*(9), 2382–2410.

Shapiro, C. and A. Yurukoglu (2024). Trends in competition in the united states: What does the evidence show? Technical report, National Bureau of Economic Research.

Sutton, J. (1991). *Sunk Costs and Market Structure: Price Competition, Advertising, and the Evolution of Concentration*. MIT Press.

Syverson, C. (2019). Macroeconomics and market power: Context, implications, and open questions. *Journal of Economic Perspectives 33*(3), 23–43.

Tremblay, V. J., N. Iwasaki, and C. H.Tremblay (2005). The dynamics of industry concentration for u.s. micro and macro brewers. *Review of Industrial Organization 26*, 307–324.

Valmari, N. (2023). Estimating production functions of multiproduct firms. *Review of Economic Studies 90*(6), 3315–3342.

Verboven, F. (2012). Empirical industrial organization and competition policy. *mimeo KU Leuven*.

A Data Appendix

In this appendix we describe the various datasets we rely on; i.e., the three demand datasets (IRI, Nielsen and Dominick's Finer Foods), the Census production data and the Compustat data. We also present some data from Consumer Reports in 1996, which provides some external information about beer markups.

A.1 External info: Consumer Reports data

Table A1: Decomposing retail prices (1996)

	Mass Brewers		Craft Beer		Overall	
Component	Item	cum	Item	cum	Overall	cum.
Retail	0.80	4.01	1.29	6.45	1.05	5.23
Distributor	0.66	3.21	1.19	5.16	0.93	4.19
Tax/shipping	0.69	2.55	0.62	3.97	0.66	3.26
Brewer profit	0.24	1.86	0.67	3.35	0.46	2.61
adv/manag	0.33	1.62	0.54	2.68	0.44	2.15
labor	0.47	1.29	1.06	2.14	0.77	1.72
packaging	0.66	0.82	0.83	1.08	0.75	0.95
ingredients	0.16	0.16	0.25	0.25	0.21	0.21
			Markups			
Brewer-low(*)	1.44		1.57		1.52	
Brewer-high(**)	2.27		3.10		2.74	
Distributor	1.26		1.30		1.28	
Retailer	1.25		1.25		1.25	

Notes: Source: Own calculations using Consumer Reports, 1996, (*) Excluding advertizing and management cost (i.e., treated as fixed costs). (**): Labor, advertizing and management costs treated as fixed costs.

A.2 Demand data

IRI and Nielsen We follow Miller and Weinberg (2017) for the aggregation of the data, and the definition of local markets and products. The empirical analysis focuses on 39 geographic regions in the IRI database. The distinct geographic markets are: Atlanta, Birmingham/Montgomery, Boston, Buffalo/Rochester, Charlotte, Chicago, Cleveland, Dallas, Des Moines, Detroit, Grand Rapids, Green Bay, Hartford, Houston, Indianapolis, Knoxville, Los Angeles, Milwaukee, Mississippi, New Orleans, New York, Omaha, Peoria/Springfield, Phoenix, Portland in Oregon, Raleigh/Durham, Richmond/Norfolk, Roanoke, Sacramento, San Diego, San Francisco, Seattle/Tacoma, South Carolina, Spokane, St. Louis, Syracuse, Toledo, Washington D.C., West Texas/New Mexico. We use the counties that compose these regions to create the equivalent markets in the Nielsen data. The potential market of each region is computed by assuming that the market size is 50% greater than the maximum unit sales within each region. As to the products, the analysis focuses on 13 flagship brands and three package sizes (6, 12, and 24/30 packs).[61] Hence, the product is defined as the combination of brand and size (e.g., Budweiser 6 pack). The market shares are measured in terms of 144-ounce equivalent units, and accordingly, prices are measured in USD per 144 ounces. Following MW, prices are converted to 2010 dollars using the CPI (All Urban Consumers) from the US Bureau of Labor Statistics. We aggregate the scanner data from the original UPC-store-week level to the product-region-month level.

We supplement the retail scanner data with two datasets. First, the income draws come from the Public Use Microdata Sample (PUMS) of the American Community Survey. Incomes were also converted to 2010 dollars using the CPI. Second, for identificaiton of demand, we use MW's data on driving miles between each market and the nearest brewery and information on diesel fuel prices. For more details about the variables and data aggregation, see MW and their supplementary material.

Below we present three tables. First, we provide an overall comparison of the two main demand datasets used in the analysis in the main text, where we compare the number of stores, markets, brands and products over time. Second, we present average prices for the leading brands, and finally we report revenue shares by brewer (which are used to aggregate up the micro markup estimates).

Table A2 indicates that Neilsen's sample of store is quite different in 2017. Some of our figures dropped this year from the time series; we note when this is the case.

[61] The flagship brands are the following: Bud Light, Budweiser, Michelob, Michelob Light, Miller Lite, Miller Genuine Draft, Miller High Life, Coors Light, Coors, Corona Extra, Corona Extra Light, Heineken, Heineken Light.

Table A2: Overall comparison of the datasets

	IRI Dataset					NIELSEN Dataset				
Year	Stores	Mkts	Brnds	Pdts	Units	Stores	Mkts	Brnds	Pdts	Units
2001	1185	39	12	35	22.0					
2002	1217	39	12	36	21.6					
2003	1222	39	12	36	23.1					
2004	1222	39	12	36	23.3					
2005	1256	39	13	38	21.8					
2006	1213	39	13	39	20.7	4689	38	13	39	102.9
2007	1211	39	13	39	20.9	5013	39	13	39	111.5
2008	1181	39	13	39	21.4	5067	39	13	39	111.7
2009	1157	39	13	39	21.6	5156	39	13	39	114.7
2010	1155	39	13	39	19.4	5175	39	13	39	108.8
2011	1091	39	13	39	17.1	5177	39	13	39	102.5
2012						5159	39	13	39	94.9
2013						5209	39	13	39	91.5
2014						4929	39	13	39	90.3
2015						4870	39	13	39	90.0
2016						4756	39	13	39	91.1
2017						2838	33	13	39	51.2
2018						5847	39	13	39	102.2
2019						5885	39	13	39	98.5

Notes: The units sold are in millions. The number of stores corresponds only to the stores selling at least one unit of beer. The segment of the stores is supermarkets (IRI) and food (Nielsen).

Table A3: Average price by brands

	IRI Dataset			NIELSEN Dataset		
	6p	12p	24/30p	6p	12p	24/30p
Budweiser	11.7	9.8	8.1	11.9	9.7	7.7
Bud Light	11.7	9.9	8.1	11.9	9.7	7.7
Coors	11.9	9.9	8.1	12.1	9.9	7.7
Coors Light	11.8	10.0	8.1	12.1	9.8	7.7
Corona Extra	16.3	14.0	12.2	16.2	13.2	11.9
Corona Light	16.3	13.9	12.3	16.2	13.1	11.8
Heineken	16.7	14.2	11.9	16.4	13.1	11.7
Heinehen Light	16.6	14.0	11.8	16.7	13.5	12.5
Michelob	12.4	10.9	8.2	12.2	10.5	7.8
Michelob Light	12.0	10.6	7.7	13.0	11.0	8.8
Miller Genuine Draft	11.9	9.7	7.9	11.7	9.2	7.5
Miller High Life	8.6	7.3	6.2	8.8	7.2	6.0
Miller Lite	11.5	9.6	8.0	11.6	9.2	7.6

Notes: Weighted (units) average price measured in equivalent unit basis of 12 pack. The calculations are based on the period 2005-2011 for the IRI database and on the period 2006-2019 for the Nielsen scanner data. Prices are reported in 2010 dollars.

Table A4: Revenue share (inside goods) by brewery

	Revenue share (inside goods)					
	IRI Dataset			NIELSEN Dataset		
	2007	2009	2011	2007	2009	2011
Anheuser-Busch InBev	39.2	39.7	40.5	32.7	33.5	33.6
Grupo Modelo	18.0	16.5	15.5	19.3	16.9	15.7
Heineken USA	9.2	9.0	8.6	10.2	9.7	9.1
MillerCoors	-	34.8	35.3	-	35.8	36.8
Molson Coors	12.7	-	-	12.7	-	-
Sabmiller	21.0	-	-	21.3	-	-
Total Revenue ($US)	204.1	208.3	160.7	1093.3	1100.1	956.1
HHI (Revenue)	2542	3140	3207	2177	2801	2838

Notes: Shares are computed using only the inside goods (see data description). Total revenue is reported in millions of 2010 dollars.

A.3 Production data

The Census of Manufactures sends a questionnaire to all manufacturing plants in the United States with more than 5 employees every five years. Labor inputs are measured by total production hours. Materials inputs are measured by total expenditure on materials including grains, malt, packaging, electricity, and fuel. Capital is measured using a question on total assets. Capital, materials and sales were deflated using the NBER-CES industry-level deflators into 1997 dollars. Our output measure is the total value of product shipments (deflated).

We construct our panel of establishments (plants) as follows. First, we collect all years of data for any establishment listed as having a primary or secondary industry affiliation in NAICS 312120 or SIC 2082. Then, we look at the number of years where that establishment has "good" data. An establishment-year qualifies as "good" if all of the following criteria are met: at least three employees, at least $50,000 in sales, materials and capital inputs are non-missing and non-zero, and the observation was included in tabulations of official census statistics. An establishment is included in our panel only if it has at least two consecutive census years of good data, and, over the sample, at least 50% of is sales come from sales of beer.

B Demand System Details

In this section, we review some important details of Miller and Weinberg's demand model and explain our alternative estimation strategy.

MW's main specification involves twelve instruments, each of which is interacted with the product-market demand shock ξ_{jdt} to form a moment:

1. An indicator for Anheuser-Busch InBev and MillerCoors products following the MillerCoors merger
2. Distance from the market to the product manufacturer's closest brewery
3. Sum of above distances for all products in market (d, t)
4. Number of products in market (d, t)
5. Interaction of 1 and 3
6. Interaction of 1 and 4
7. Indicator for Anheuser-Busch products interacted with 3
8. Indicator for Anheuser-Busch products interacted with 4
9. Mean income in market (d, t)
10. An indicator for imported products interacted with 9
11. Calories interacted with 9
12. Package size interacted with 9

MW also include fixed effects for each product and time period. We follow MW in dropping the merger period of June 2008 through May 2009, so the post-merger indicator is equal to 1 for June 2009 onward for products owned by either Anheuser-Busch InBev or MillerCoors.

When we drop the cost-shifters from estimation, we drop the moments associated with 2, 3, 5, and 7. When we drop the BLP instruments, we drop moments 4, 6, and 8.

Table 4 was generated using a modified version of MW's code. Because we want to illustrate the difference that adding and/or dropping moments makes, rather than presenting differences resulting from changes in how the moments are weighted, we avoid changing how moments are weighted from column to column. To be precise, column 1 is estimated using standard two-step GMM estimation, following MW. For the other columns, we start with the same weighting matrix used in the second step for column 1's (MW's) estimation. To produce columns 4 and 5, we simply drop the rows and columns corresponding to the dropped moments, and then we run the second-step GMM estimation. For columns 2 and 3, we do the same thing, but we also add another row and column to the weighting matrix for the production-based markup moment. The elements of the new row and column are set to zero, except for the diagonal element, which is set equal to the inverse of the variance of the fitted values of the new moment, evaluated at the first-step parameter estimates. That is, the weighting matrix is given by

$$\begin{bmatrix} W^{MW} & \mathbf{0} \\ \mathbf{0}' & Var\left(\mu_{jdt}^{MW} - \hat{\mu}\right)^{-1} \end{bmatrix},$$

where W^{MW} is the second-stage weighting matrix from MW's main specification, selecting the rows and columns that correspond to moments we are using in our estimation, and $Var\left(\mu_{jdt}^{MW} - \hat{\mu}\right)^{-1}$ is the variance of our markup moment, evaluated at MW's first-stage parameter estimates.

The addition of the markup moment calls for a subtle but important change to the typical strategy for implementing differentiated products demand estimation following Berry et al. (1995). Demand parameters are typically divided into "linear" and "nonlinear" parameters, where the nonlinear parameters must be explicitly searched over when minimizing the GMM objective function, and the linear parameters can be concentrated out, or solved for algebraically given the candidate nonlinear parameters—see section 6.5 of Berry et al. (1995). Normally, parameters controlling heterogeneity in price sensitivity are nonlinear parameters, and the mean price sensitivity parameter is a linear parameter. However, with the inclusion of the markup moment, the linear price parameter must be treated as a nonlinear parameter because it affects the markup moment, which means we no longer have linear first-order conditions for the optimal value of the mean price parameter.

A final detail is accounting for uncertainty in the production-based moment when computing standard errors for the demand estimates. We do this by simulating the asymptotic distribution of the demand estimates, using the asymptotic distribution of the production-based markups. That is, we start by simulating the asymptotic distribution of production-based markup estimates, drawing 100 points from the estimated distribution. For each of these draws of a production-based markup, we implement our modified estimation of MW's model imputing the drawn markup value in the new moment. The standard errors in column 2 of Table 4 are then computed by integrating over these 100 points and the standard GMM asymptotic distribution associated with the demand estimates for each draw.

C Additional Results

C.1 Markups Time Series and the Role of Assumptions

Figure C1: Demand-Based Markups with Alternative Assumptions

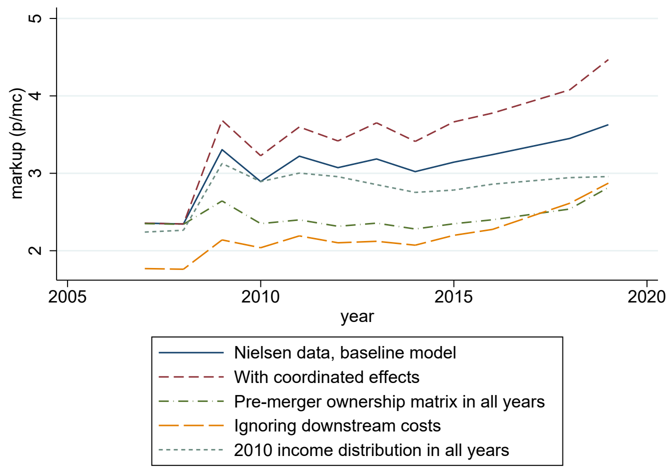

Notes: We plot the sales-weighted average demand-based markup using the Nielsen data. Following MW, the merger period of June 2008 to May 2009 is omitted from these calculations. Following MW, the merger period of June 2008 to May 2009 is omitted from these calculations. 2017 was dropped because the sample of stores was very different that year.

Focusing on the Nielsen data, Figure C1 shows how the demand-based markups time series depends on modeling assumptions. First, following MW's results, we assume that after the Miller-Coors merger, there were coordinated effects such that Anheuser-Busch and Miller-Coors internalized each other's profits, valuing each dollar of the other firm's profits as .264 dollars of their own profits. In this case ("With coordinated effects"), we see a somewhat larger increase in markups following the 2008 merger. Going in the opposite direction, we can pretend the merger never happened, and use the pre-merger ownership matrix throughout the sample. In this case, markups after 2008 are substantially lower.

Figure C1 also makes it clear that accounting for downstream costs (retail and distribution costs and taxes) makes a substantial difference in the level of markups. Finally, we compute a time series in which we use the 2010 income distribution for all years, shutting down changes in markups that are driven by demographic change and the associated changes in demand parameters. In this case, markups are largely flat after the merger, and the merger is the only substantial reason that markups increase during this sample.

Another feature of our baseline markup calculations is that we top-code markups at 15 and replace negative markups with a markup of 15. The motivation for doing this comes from the fact that, as a single-product firms' own-price elasticity of demand approaches -1 from below, the markup asymptotes to infinity. However, an own-price elasticity of slightly larger than -1 (or smaller than 1 in absolute value) implies negative markups. This discontinuity in the mapping from demand elasticities to markups can lead to strange comparative statics, in the sense that something that makes demand less elastic will tend to increase markups, except when elasticities cross a threshold that discontinuously makes the markups jump from very large positive numbers to negative numbers. Our top-coding strategy avoids such discontinuities.

Figure C2 shows the markup time series without top-coding. While the topcoding makes only a small difference in the time series based on the IRI data, the Nielsen-based series in 2011 illustrates the possibility that results can be implausible and unstable without top-coding. Average markup levels are sensitive to the possibility that a single product can have an astronomical markup (e.g., in the case of a single-product firm, when own-price demand elasticity is slightly larger than 1 in absolute value).

Figure C2: Demand-Based Markups with Alternative Assumptions

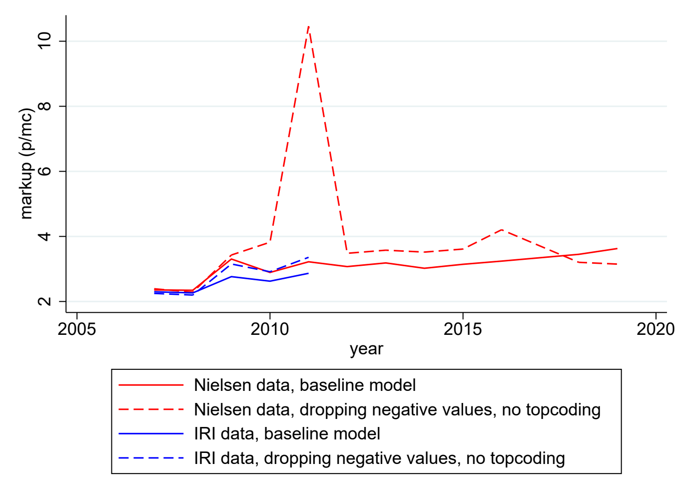

Notes: We plot the sales-weighted average demand-based markup using the IRI and the Nielsen data. The baseline calculations impose a maximum markup of 15 and sets all negative markups to 15. The versions without topcoding do not impose a maximum, and negative markups are dropped. Following MW, the merger period of June 2008 to May 2009 is omitted from these calculations.

C.2 Retail Competition Without Costs of Retailing

Our main results in Section 5.2.3 imputes (small) retail costs as described in Section 3.1.2. Figure C3 adds a specification with zero cost of retailing.

Figure C3: Markups, retail market structure, and costs of retailing

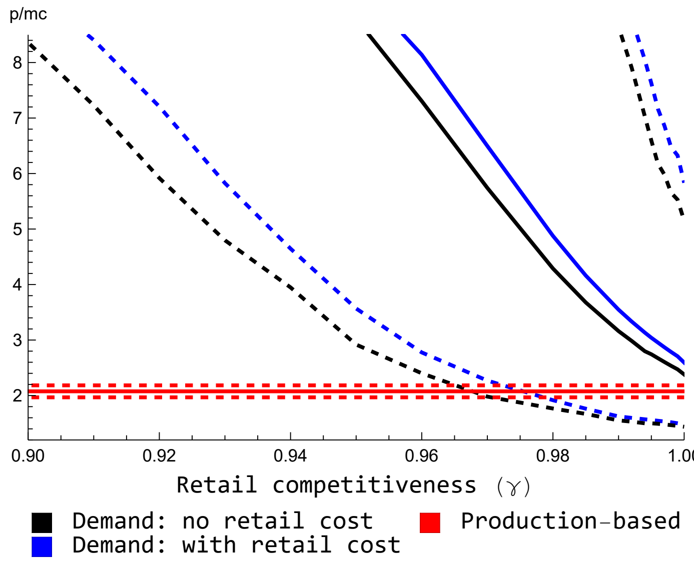

Notes: The blue lines represent implied markups for values of the retail competitiveness parameter (γ), and the dashed lines indicate the 95% confidence interval, with costs of retailing imputed as described in the text. Black lines display the analogous demand-based results imputing a cost of retailing of zero. The solid red line represents the estimated (average) markup using the production approach (using the Leontief-Cobb-Douglas specification), while the dashed red lines represent its 95% confidence interval.

C.3 Demand-based markups with alternative models

Our earlier draft estimated demand using the Dominick's Finer Foods dataset, which covers an earlier time period (1992-1995). These data only covers the Chicago area, and comes from one retailer, but these data has been used extensively to study the beer market and it has a few important advantages, chief among them reporting both retail and wholesale prices. For details about the data, model, and estimation, see the subsumed draft `https://www.nber.org/papers/w22957`. Figure C4 adds the mean demand-based markups from that draft (in 1992, which corresponds to a Census year on the production side) to our markups time series.

Our current draft uses the Dominick's data only to measure the average gap between retail and wholesale prices as a measure of the costs of retailing, as described in section

3.1.2.

Figure C4: Aggregate markups across Methods and Datasets

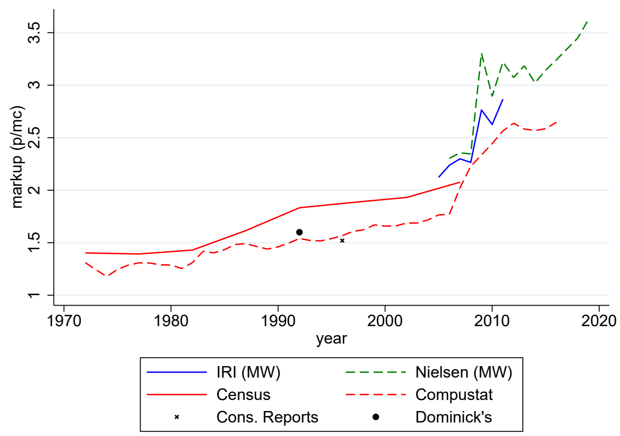

Notes We plot the sales-weighted average markup across firms (using plants, product-market micro estimates of markups) using the production approach (red lines) on the census and Compustat data, and using the demand approach with either the IRI or Nielsen data. The solid purple line presents markups from Hildalgo's (2023) alternative demand specification. The black dot are the demand-based markups from our earlier draft, which used a different demand specification and Dominick's data. We also include insert markups inferred from an article in Consumer Reports (1996). All demand-based markups are computed assuming the benchmark passive retailing assumption, where we correct for the downstream retail and wholesale costs as described in the text. Following MW, the merger period of June 2008 to May 2009 is omitted from the demand-based calculations. 2017 was dropped from the Nielsen time series because the sample of stores was very different that year.

Comparing markups across the two approaches for the year 1992 yields a very similar conclusion to what was presented in the main text for 2007. Table C1 reports, for each approach, the implied markups for the years where we have census data as well as demand estimates.

Hidalgo (2023) features a beer demand model with richer consumer heterogeneity. The demographics include income, a dummy variable indicating whether the individual belongs to either the Millennial or Gen Z generation, and a dummy variable indicating whether the individual self-identifies as being of Hispanic ethnicity. Figure C4 plot Hidalgo's markups assuming Nash-Bertrand competition and using a correction for downstream costs similar to ours. For further details of the dataset, demand model, and identification, see Hidalgo (2023), particularly in Sections 2.4, 4.1.2, and 5.1.

Table C1: Markup Estimates Across Approaches, Data Sources, and Years

Year	External	Production		Demand		
	Cons. Rep.	Census	Compustat	Dominick's	IRI	Nielsen
1992		1.68	1.54	1.65	-	-
1996	1.52	-	1.57	-	-	-
2007		2.08	2.02	-	2.29	2.18

C.4 Production function parameters

Table C2: Production function parameter estimates

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
betal	0.125	0.298	0.436	1.390	0.538	1.615	0.749	0.766
	(0.110)	(0.185)	(0.075)	(0.216)	(0.058)	(0.241)	(0.118)	(0.359)
betam	0.700	0.717						
	(0.148)	(0.248)						
betak	0.177	0.164	0.538	-0.152	0.489	-0.106	0.300	0.235
	(0.061)	(0.124)	(0.098)	(0.196)	(0.065)	(0.200)	(0.123)	(0.259)
betal2		0.042		0.127		0.168		0.271
		(0.048)		(0.028)		(0.027)		(0.272)
betam2		0.054						
		(0.023)						
betak2		0.008		0.094		0.102		0.084
		(0.025)		(0.019)		(0.019)		(0.076)
betalm		-0.124						
		(0.056)						
betamk		-0.048						
		(0.046)						
betalk		0.052		-0.230		-0.279		-0.287
		(0.044)		(0.042)		(0.045)		(0.267)
betalmk		0.001						
		(0.002)						
Specification	GO	GO	VA	VA	LEO	LEO	LEO	LEO
Technology	CD	TL	CD	TL	CD	TL	CD	TL
Labor IV	l_t	l_t	l_t	l_t	l_t	l_t	l_{t-1}	l_{t-1}

Notes: Standard errors in parentheses. Technology is either Cobb-Douglas or translog. Specification is either gross output, value added, or Leontief. Instruments include k_t and either l_t or l_{t-1}. Instruments also include m_{t-1} for gross output specifications and interactions for translog specifications. All specifications shown here assume that expected productivity is linear in lagged productivity.