Global giants and local stars: How changes in brand ownership affect competition*

Vanessa Alviarez[†] Keith Head[‡] Thierry Mayer[§]

June 1, 2023

Abstract

Large beer and spirits makers expanded abroad mainly by acquiring local brands. Exploiting market share data in 76 countries and changes in brand ownership from 2007 to 2018, we estimate that owners matter little for brand performance, except via a negative consumer response to foreign ownership. Our counterfactuals indicate that market power increases were large enough to yield higher profits for the majority of mergers without relying on fixed-cost savings. Emulating pro-competition policies used by the US and EU could have saved South American consumers up to 18%. US beer prices would be 3–4% higher without DOJ-enforced divestitures.

Key words: multinationals, oligopoly, markups, concentration, firm effects, brands, frictions, mergers and acquisitions, competition policy

JEL classification: F12, F23, F61, K21, L13

^{*}This research has received funding from Spain's State Research Agency MDM-2016-0684 under the María de Maeztu Programme, a SSHRC Insight grant, and the Centre for Innovative Data at UBC. We thank seminar and workshop participants at University of Texas, University of Michigan, University of Hong Kong, the Canadian Economics Association, the Rocky Mountain Empirical Trade workshop, ER-WIT, Nottingham University, CEMFI, the Western Economic Association, Université Libre de Bruxelles (ULB-ECARES), KU Leuven, Aarhus University, Copehagen Business School, the Geneva Trade and Development Workshop, Inter-American Development Bank, Carlos III, and Yale University for comments and suggestions. We thank Jan de Loecker for an insightful discussion. Raffaele Saggio, Stephane Bonhomme, Koen Jochmans guided us on econometric methods for network data. Nathan Miller and Michele Fioretti provided helpful advice and Rodrigo Heresi provided data. Behrad Hafezi supplied essential research

[†]Research Department, Inter-American Development Bank, valviarezr@iadb.org

[‡]Sauder School of Business, University of British Columbia, CEPR, CEP, CEMFI (visiting 2019–2020) keith.head@sauder.ubc.ca

[§]Sciences Po, Banque de France, CEPII, and CEPR, thierry.mayer@sciencespo.fr

1 Introduction

Concern—and controversy—over the rise of market power has spread well beyond competition policy specialists and industrial organization economists. One reason is the attention-grabbing findings of rising concentration and markups. Grullon et al. (2019) report that concentration indexes increased in three quarters of US industries from 1997 to 2014. De Loecker et al. (2020) and De Loecker and Eeckhout (2018) show rises in sales-weighted markups in the US (from 1.2 to 1.7) and globally (from 1.1 to 1.8) since 1980. Such observations have kindled debate over the mechanisms that might drive widespread increases in markups. Reviewing other major phenomena documented during the same period (1980–2016), it is natural to ask what role globalization might play. Intuitively, lower trade and investment frictions should increase competitive pressure and thereby *decrease* markups.¹ However, this reasoning ignores a number of mechanisms that could push markups in the opposite direction.

There are at least three channels through which globalization might increase markups. Recent research has investigated two of them. Autor et al. (2020) propose that "greater product market competition (e.g., through globalization)" has allowed the most productive firms—with the highest markups—to increase their market shares. Thus, *aggregate* (share-weighted) markups can rise even in an increasingly competitive world.² A very different channel works through imported inputs: decreases in input tariffs tend to lower the overall costs of production. When firms fail to pass on those cost reductions completely, markups rise (De Loecker et al., 2016).³ A third mechanism for globalization to raise markups is via growth in cross-border mergers and acquisitions (M&A). As large multinational corporations (MNCs) absorb previously competing entities, the acquiring firms have the ability and the incentive to increase markups.

This paper focuses on this third channel, estimating and quantifying the ways that ownership changes affect competition in two beverage industries, beer and spirits. A key to understanding the market power effect of international mergers is found in the market interactions between brands referred to as "global giants" and "local stars." The former are MNC-owned brands sold in many countries, whereas the latter brands gar-

¹Brander and Krugman (1983) is a pioneering model of the "pro-competitive" effects of trade liberalization in which markups fall along with lower transport costs.

²Autor et al. (2020) marshal evidence supporting a rise in aggregate markups through what they call the "superstar firm framework." (Syverson, 2019, p. 27) and (Berry et al., 2019, p. 58) develop variations on this composition argument.

³This paper finds that Indian tariff reductions led to rising markups through this channel. The World Bank (2020) reports that global value chain participation has increased markups of large corporations in developed countries.

ner high market shares only in their country of origin. AB InBev's purchase of the beer brands owned by the Grupo Leon Jimenes provides a useful example. Before the sale, Leon Jimenes' Presidente and Presidente Light dominated the lager and light beer segments in the Dominican Republic. AB InBev had smaller shares of these segments with its brands Brahma and Brahma Light. The combination of global giant brands with local stars, leading to a 98% share of beer sold in the Dominican Republic motivates AB InBev to elevate and harmonize brand-level markups.

Not all governments were passive during the recent phase of multinational brand amalgamation. The US and EU authorities in particular intervened to force acquiring firms to divest brands in markets where they deemed the mergers to have anti-competitive effects. For example, AB InBev had to transfer the US market rights on Corona to Constellation Brands when it acquired the parent company, Grupo Modelo. Later, the EU compelled AB InBev to divest Peroni and several other European brands to Asahi after the acquisition of SABMiller in 2016. This form of "structural remedy" is attractive because it dis-incentivizes firms from raising markups. However, the potential downside to forcing divestitures is foregone efficiencies. For example, AB InBev claimed its 2008 purchase of Anheuser Busch had generated \$2.3bn in annual savings and that buying Grupo Modelo would lead to a further \$600mn per year.⁴ The need to quantify the consequences of divestitures motivates this paper's estimates of how new ownership affects the costs and appeal of the acquired brands. We conduct counterfactuals applying these estimates within a multi-product oligopoly model, considering the impact of more and less permissive mergers policies on a price index computed in a way that includes changes in appeal.

This paper centers around two distinct empirical exercises. In the first, we estimate changes in the cost-adjusted appeal of a brand following acquisition by a new owner, often headquartered in a different country. The second exercise plugs those estimates into a calibrated oligopoly model to solve for new equilibrium prices in each country impacted by mergers. In both exercises, we use a three-layer nested demand structure and static Nash markup determination.⁵

Our decision to estimate appeal/cost changes, but use counterfactuals to compute price changes is based on the relative strengths of our data set and our view of the most important knowledge gaps in the literature. The empirical pattern of prices following

⁴*Financial Times*, "AB InBev/Modelo: no cheap round" June 29, 2012.

⁵Pinkse and Slade (2004) find that static Nash oligopoly in prices is not rejected in the British beer market. Miller et al. (2021) argue that conduct in the US beer industry is better characterized by price leadership. This conduct exacerbates the price-increasing effects of mergers as compared to Bertrand. Throughout this paper we consider both Bertrand and the "softer" competition implied by Cournot conduct.

mergers have been investigated in surveys of 49 studies by Ashenfelter et al. (2014) and 47 studies by Kwoka (2014). The results mainly (36 out of 49 in the former survey, 73 of 119 mergers in the latter) support the oligopoly prediction that merger-driven concentration increases lead to higher prices. Ashenfelter and Hosken (2010) find significant price increases ("typically between 3 and 7 percent") in four of five mergers they study, including one very relevant for this paper, the merger that created Diageo. Dafny et al. (2012) established the methodology of regressing change in log price on the change in concentration predicted by a naive merger analysis. They report significant causal effects of merger-induced concentration on premiums in the insurance industry. Ashenfelter et al. (2015) and Miller and Weinberg (2017) estimate similar regressions exploiting geographic variation within the US to show that merger shocks to the Herfindahl concentration index increase the price of beer.

The mechanism linking mergers, rising concentration, and price increases thus receives firm empirical backing from high-quality studies of multiple sectors. However, this body of work mainly considers the US market.⁶ Since many of the largest mergers involve cross-border acquisitions, there are two important knowledge gaps. First, how do the consequences of multinational mergers vary across affected countries depending on their initial market structures? Second, are consumers helped or harmed when foreign multinationals acquire their favored local brands? The data we employ are uniquely well qualified for these tasks as they track brand ownership and market shares for all major markets (76 countries in all) during a decade featuring widespread ownership changes. As initial market structures and policy permissiveness vary widely, some countries experience much larger increases in concentration than others. In particular we find that beer and spirits mergers tended to redistribute surplus from poorer countries to the high income countries where the multinationals are headquartered.

The core quantitative analysis in this paper computes markups under the observed set of ownership relationships before comparing those markups to those that would have arisen in alternative scenarios. There are two prominent methods of revealing markups. The first method, pioneered by Berry (1994), relies on the first-order conditions linking marginal revenue to marginal cost under particular conduct assumptions. Once a demand curve has been estimated, the ratio of price to marginal cost can be inferred. A second markup method, developed by De Loecker and Warzynski (2012), eschews conduct assumptions and instead reveals markups from the firms' cost minimization problem. It relies on input use data and estimated production function parameters. We follow the first approach here for three reasons. First, we lack data on firm-level input use that is

⁶Kwoka (2014) restricts attention to mergers that affected the United States.

critical for the production function approach. Second, even if we could observe input use for all the firms in our data set, one cannot use the production function approach to determine markups in different countries without imposing additional structure to allocate input use across markets.⁷ Third, and most importantly for our purposes, the structure imposed in the demand-side method is well-suited to computing markup changes in response to counterfactual reallocations of brands to different owners.

Our analysis combines tools from industrial organization, labor, and trade. Each has precedents in their own field but here we show how they each contribute to understanding the impact of cross-border mergers. From IO we have a demand structure incorporating richer substitution patterns than are customary in trade. Most features in our model come from the constant expenditure three-layer nested demand developed by Björnerstedt and Verboven (2016). Our inference of cost-adjusted appeal draws on the Hottman et al. (2016) method for backing out product appeal and the Nocke and Schutz (2018) idea of the products having type that is a sufficient statistic for market share. The key features are multi-product oligopoly and nested constant elasticity of substitution (CES) demand.⁸ We also employ the multi-product oligopoly structure that is standard in IO. These tools allow us to obtain a more trustworthy calculation of how bringing diverse brands under the same ownership will affect markups. From trade, we use a brand-market level gravity equation to estimate home bias as well as related effects of proximity and shared language on the pattern of market shares of beverage brands in different countries. From labor, we adopt techniques for measuring employer contribution and sorting but apply them to brand/owner instead of employee/employer. One of our key finding echoes a finding in the employer/employee studies: after correcting for bias, firms' contributions to outcomes are much smaller than the shares implied by naive fixed effect methods.

Our paper contributes four key findings. First, we quantify across all major markets the potential savings to consumers from forcing divestitures of brands as a condition of merger approval. Relative to the counterfactual of a permissive merger policy, the actual remedies imposed on AB InBev lower the price index for US beer by three to four percent, depending on the conduct assumption. Conversely, passive countries paid as much as 18% more for beer than they would have by emulating US and EU remedies. Our second contribution is to show that the specific owner of a brand contributes surprisingly little to its performance. After mitigating the upward bias caused by limited brand mobility, firm

⁷De Loecker et al. (2016) devise an input allocation method for firms that sell multiple products.

⁸Departing from the pioneering work in trade using nested CES structures, Atkeson and Burstein (2008) and Edmond et al. (2015), our model add layer between brand and industry. Furthermore, brands are allowed to be large at all three levels.

effects explain 4% or less of the variation in a brand cost-adjusted appeal. Simulations show that correcting for endogenous mobility would further lower the owner contribution. This result is reassuring for antitrust policy since it implies that forcing a brand divestiture will probably not sacrifice large owner-derived benefits. A third important result is that the *geography* of ownership matters. An owner with a foreign headquarters tends to lower cost-adjusted appeal in a market by 11–12 percent. We believe this is the first study to estimate this type of negative impact of overseas ownership.

In addition to the substantive findings described above, our paper makes two methodological advances. Most importantly, we show how to adapt the exact hat algebra approach pioneered in Dekle et al. (2008) to run counterfactuals in settings where a few large multi-product firms interact as oligopolists, while a fringe of individually small firms price as in monopolistic competition. This generalization is valuable because it offers a framework for addressing oligopoly issues that is more economical in its data requirements than the standard industrial organization approach. Second, we show how to apply recent techniques from labor economics to diagnose limited mobility bias and mitigate its impact on the estimated contribution of firms.⁹ This application in the context of measuring owner value-added in product markets provides a template for research on related questions.

The remainder of the paper proceeds as follows. Section 2 describes the data we use, highlighting its advantages and limitations. Section 3 presents the model. There we describe the method to back out cost-adjusted appeal for each brand in each market. Section 4 estimates the effects of firm ownership on this determinant of brand performance. Here we exploit the cross-country variation in our data that permits estimation with brand-firm interactive effects. Using estimates of the systematic changes in cost-adjusted appeal associated with the identity and headquarters of the owner, we compute counterfactuals in section 5 for alternative patterns of ownership that might have prevailed in 2018 had different merger policies been adopted. To assess the combined impact of a decade of mergers on prices, concentration, and profits, we also consider the counterfactual of restoring every brand to its 2007 owner.

2 Data: sources and patterns

Our dataset combines four distinct components. The first of those provides sales at the brand-country-year level. Crucially, this data tracks the ultimate owner of each brand in

⁹Jochmans and Weidner (2019) provide the diagnostic (connectivity) measure and Andrews et al. (2008), Bonhomme and Manresa (2015), and Kline et al. (2020) provide the mitigation techniques.

a given period. The second dataset, created as part of this study, determines the origin (Corona originates in Mexico, Absolut from Sweden) and module (Corona is a lager, Absolut is a vodka) of 4,352 brands. The third, also original to this study, identifies the headquarters country for each of the 1,023 firms owning beer and spirits brands. Finally, we use standard data (available from CEPII) on bilateral distances and common languages.

2.1 Market shares and ownership

The Global Market Information Dataset (GMID), from Euromonitor, reports sales information for individual brands and their corresponding owners for specific consumer products in 75 to 80 countries for the most recent 10 years. By combining two "vintages" of the data, we obtain a sales panel running from 2007 to 2018. Within each combination of product category, market, and year, GMID lists sales for all brands above a threshold market share, which the documentation lists as 0.1%. GMID sums the sales of smaller brands in a given market and lists their collective shares under the brand names "Private Label" and "Others." Private Label has less than 1% market share in the median country for both beer and spirits. The market share of Others is generally small for beer (median of 11%), but accounts for one third of the German market. In the US, Others have risen from 11% in 2007 to 20% in 2018. Liquor markets are more fragmented, with Others accounting for a median of 26% of sales. We calculate the shares of brands and firms in each national market using as a denominator the sales of all brands, including Others and Private Label, which we refer to collectively as the fringe.

GMID tracks all changes in majority ownership at the brand level occurring over the 2007–2018 period. This feature is distinctive in that most M&A datasets record changes in ownership at the firm level, without providing explicit information about which product lines or brands are involved in the transaction.

The GMID market share data addresses several concerns regarding concentration measures derived from the economic census or firm-level databases such as Compustat and Orbis. First, markets are defined from the consumer point of view, considering horizontal substitutes. Other databases rely on standard industry classifications that were mainly designed to capture similarities between firms. Berry et al. (2019) point out that "industrial classifications in the Census often fail to reflect well-defined economic markets." They give the example of software, but an example given by Grullon et al. (2019) provides a more striking illustration. One of their 3-digit NAICS industries is leather products. Sub-industries include handbags and footwear, two products we might think of as complements. Another sub-industry, leather tanning, should be thought of as an input to the other two. It makes little sense to think of a firm with a high share of aggregate production in leather products as having market power in a particular consumer market. The firms in the beverage categories we study compete with each other through their portfolios of substitute brands.

A second advantage of GMID for calculating market shares and concentration in a way that is relevant for markups is that we see brand-level sales in a given market including imported products. Other data sets such as the census or Compustat report the revenue of a set of firms, aggregating over all markets. Such revenue measures include exports to other markets, but exclude imports. Thus, census data does not measure sales in the market in question.¹⁰ Imports supplied by foreign firms should increase competition. On the other hand, imports carried out by large domestic firms, with little or no local production, can actually increase concentration relative to measures based on domestic shipments. Our data overcomes these issues since brand sales aggregate to total expenditures in a market.

Studies of concentration using Compustat omit private companies, which include a few large firms (e.g. Bacardi) and the often large fringes of small firms. Both Compustat and census omit sales of multi-category companies outside their assigned SIC. This issue could be quantitatively important since Compustat classifies Pernod Ricard, the second largest spirits distiller in the world, as a winery.

Category	Brands		Firms	Countries		ies	
	All	multiple			HQ	Origin	Market
		markets	owners				
Beer	2425	368	672	464	79	93	78
Spirits	2894	598	528	849	87	106	77

Table 1: Beer and spirits makers and their brands

Notes: The market shares and owner-brand relationships come from GMID; origins and headquarters were collected by authors.

Table 1 shows that each category comprises hundreds of firms and most categories have thousands of brands. The regression method we use to estimate firm ownership effects on brand performance depends on observing the same brand sold by different firms and in different markets. Compared to other beverages in the GMID data, beer and spirits stand out as having large numbers of brands that changed ownership. As shown in the third column, 28% of the beer brands in the data set had more than one owner. This includes a few brands, such as Corona and Fosters, that have different owners in

¹⁰Compustat has the larger concern that it mainly reports consolidated data which includes sales from majority affiliates in other countries than the one where the firm is based.

different markets. Spirits also exhibits substantial mobility of brands across owners, with about 18% having more than one owner. Spirits has the highest count of multi-market brands, which is important for backing out both brand effects and brand-origin frictions. The last three columns illustrate the diversity of headquarters countries, brand origins, and markets represented in the data.

Dividing Beer and Spirits brands into modules (the term Nielsen uses for market segments) within the major categories is important in the context of competition within these industries because substitution should be stronger between brands in the same module. For example, Leffe and Stella Artois, two brands owned by ABI, should be thought of as competing within different modules (ale and lager, respectively). The EU competition authority focused on modules when it evaluated the effects of spirits mergers. The European Commission (2008) report on the merger between Pernod Ricard and V&S spirits drew attention to the sharp increase in concentration of the gin market in Poland even though both the Pernod Ricard brand (Seagram's Gin) and the V&S brand (Gin Lubuski) had negligible shares of the overall Polish spirits market.

Table 2 shows the number of brands (*n*) and calculates five different types of market share: modules within groups $(s_{m|g})$, brands within groups $(s_{b|g})$, brands within modules $(s_{b|m})$, as well as the firm-level aggregations of the last two. The beer market in most countries is dominated by lager $(s_{m|g} = 63.7\%)$, so firms that have a high share of the lager market also have a high share of beer. Shares within the other modules are much larger than in lager. For example in light beer (including 0% alcohol), the average firm has $s_{f|g} = 17.5\%$ share of the beer market, but a $s_{f|m} = 47.9\%$ share of the light beer module. No single module dominates spirits. The large shares for liqueurs is because that is the classification given to the various national spirits (baijiu in China, shochu in Japan, anisettes in Turkey). Aside from vodka, the average firm share of a module ranges from 24%-37%.¹¹

2.2 Corporate headquarters and brand origins

GMID lists the global ultimate owner for each brand. This is based on majority ownership and omits the minority share positions that the multinationals sometimes take.¹² The headquarters country of each company in the GMID dataset is obtained by combining information from Orbis (Bureau van Dijk), the historical Directory of Corporate Affiliations from Lexis-Nexis, and Capital IQ. Matching the name of each brand's owners

¹¹The methods used to divide brands into modules are described in Appendix A.

¹²For instance, GMID lists China Resources as the owner of the Snow brand even in the years when SABMiller owned 49% of China Resources.

Module	n	$s_{m g}$	$s_{b g}$	$s_{b m}$	$s_{f g}$	$s_{f m}$		
	g = Beer							
lager	1277	63.7	2.8	4.4	19.6	23.2		
ale	231	9.2	0.8	20.7	13.9	33.8		
pilsner	208	16.3	3.3	25.8	22.3	42.5		
light beer	177	12.9	2.3	33.0	27.5	47.9		
malt	76	4.5	1.3	51.5	28.6	67.3		
wheat	62	3.7	0.8	44.8	12.8	54.4		
stout	34	3.8	1.4	63.7	21.1	69.7		
porter	33	1.6	0.7	55.7	19.1	70.8		
-			g = S	Spirits				
liqueurs	945	21.1	0.9	5.6	7.0	13.3		
vodka	478	17.6	1.2	9.7	8.7	18.1		
rum	313	10.8	1.1	13.2	9.3	23.5		
brandy	269	6.5	1.0	21.4	5.9	30.4		
whiskey	242	11.3	1.0	14.8	8.2	25.0		
scotch	215	17.8	1.1	8.5	10.6	29.8		
gin	174	5.4	0.8	20.5	11.9	32.0		
tequila	97	7.3	0.6	28.4	4.6	36.3		
cognac	95	6.5	0.9	26.8	7.9	36.6		

Table 2: Descriptive statistics on modules within beer and spirits

Note: All market shares are averages across brands, markets, and years.

in the GMID dataset with the names of firms in those datasets, we take the headquarters to be the location of the firm highest up the hierarchy of ownership. The exceptions are where this ultimate owner appears to be a holding company located in a tax haven. In those cases, we do additional investigation to assign a HQ location that corresponds to the place where management decisions are taken.

In one important case, AB InBev, we consider the firm to have dual headquarters, the US and Belgium. While the official head office remains in Belgium, New York City is listed as a second "Global Headquarters" on the www.ab-inbev.com site. According to reporting in the St. Louis Post-Dispatch (15 July 2018), "many key corporate functions, including a bulk of sales and marketing positions, now operate out of New York City." We set the headquarters as varying by market depending on whether the US or Belgium is closer and treat the firm as having three official languages (English, French, and Dutch).

The origin of a brand is the country where it was developed and introduced. Thus Lagunitas is an American brand and Tecate is a Mexican brand even though both are currently owned by the Dutch firm Heineken NV. Generally speaking, the origin coincides with the country where an independent firm founded the brand. We determined origins for brands by combining information from crowd-sourced product rating websites (e.g. ratebeer.com), Google Images, corporate websites, news articles, Wikipedia, and trademark registries.

2.3 Visualizing multinational brand amalgamation

Figure 1 and 2 illustrate the rise in market shares, brand ownership, and diversity of brand origins for the seven largest companies in the beer and spirits industries. The left panel of each figure shows the growth of market share. AB InBev goes from 11% to 26% of the world beer market.¹³ Heineken, Asahi, and MolsonCoors for beer, Diageo and Suntory for spirits also register visible gains. The center plot shows that these firms have even more notable increases in the number of brands. The right panel of each figure shows that, by 2018, the top beer makers had brands from around 40 countries in their portfolios. The top spirits makers held brands from about 25 brand origins each (though Pernod Ricard appeared to be retreating from international diversification).

Figure 1 shows that the most active firm in mergers and acquisitions during the period we study was AB InBev. In our model, the rising firm-level market share maps into rising markups. Therefore markups within the model follow a similarly impressive increase. It is natural to ask whether markups in the accounting data rose during the same period.

¹³InBev (11% market share in 2007) merged in 2008 with Anheuser Busch (8%) to form AB InBev.

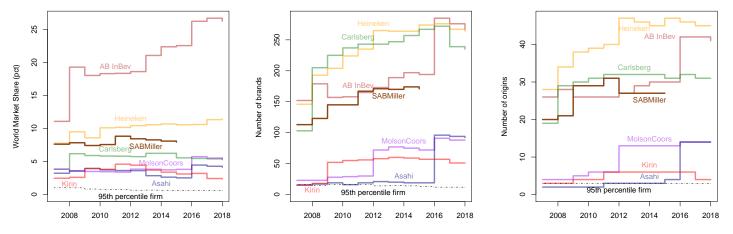


Figure 1: The growth of beer multinationals

Notes: In 2008 InBev purchases Anheuser-Busch and Heineken and Carlsberg jointly purchase Scottish & Newcastle (along with BBH) and redistribute the acquired brands among themselves. In 2009 AB InBev sells off Korean and East European brands (forming Starbev) and Kirin acquires Lion (NZ). In 2012 MolsonCoors buys Starbev and Heineken buys Asia Pacific Breweries. In 2016, AB InBev buys SABMiller, while divesting some SABMiller brands to MolsonCoors and others to Asahi to comply with antitrust orders.

30 250 Pernod-Ricard 9 25 200 Diage ¢ World Market Share (pct) Diageo Diage 20 Pernod-Ricard Number of origins Number of brands 150 Suntor 9 Bacardi 15 100 4 Suntory 9 Suntor Bacard Brown-Formar Bacard 50 2 ŝ -Formar 95th percentile firm 95th percentile firm 0 0 0 95th percentile firm 2008 2010 2012 2014 2016 2018 2008 2010 2012 2016 2018 2008 2010 2014 2016 2018 2014 2012

Figure 2: The growth of spirits multinationals

Notes: In 2008 Pernod-Ricard buys Vin & Sprit (owner of Absolut and 74 other brands). In 2014 Suntory buys Beam (which had been spun off from Fortune Brands in 2011) and Diageo buys UB Group.

Following the methodology of De Loecker et al. (2020), we calculated the variable profits and profits net of capital expenditures of AB InBev from 2000 to 2019. We then divided the profit measures by sales to obtain measures of profitability. The variable (blue) and net (magenta) profitability series are plotted in Figure 3. The Belgian-based Interbrew completed three major mergers during these two decades to yield the current AB InBev. The year of each merger is depicted with a black dotted line and we shade the subsequent two years in gray. We see that profitability rose in each of these periods. Indeed, it appears that almost all of the overall increase in profitability occurred during the adjustment period after each of those mega-mergers. The fact that profits net of capital expenditures (rK) rose supports the interpretation that AB InBev's brand amalgamation process led to rising market power.

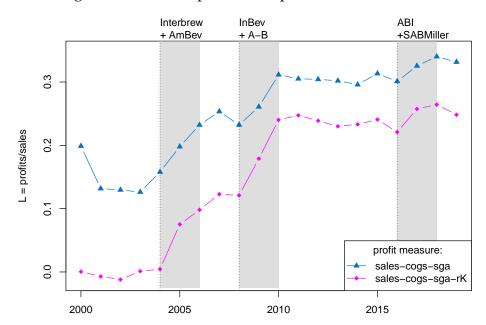


Figure 3: As ABI expanded, its profit/sales ratios rose

Source: Income statement and balance-sheet data from Refinitiv (2022). The cost components are cogs (cost of goods sold) and sga (selling, general and administrative), and rK (capital expenditures measures as gross property, plant and equipment multiplied by the user cost, i.e. interest rate minus inflation plus a 12% depreciation rate). See De Loecker et al. (2020) replication code at https://doi.org/10.7910/DVN/5GH8XO

In table 3, we turn to the case of Diageo, the largest and most multinational of spirits makers, which was formed in 1997 as a merger of Grand Metropolitan and Guinness. It dramatically expanded its portfolio of spirits brands when taking over the brands of the failing Seagram company in 2001. On its website, Diageo distinguishes between its portfolios of "Global Giants" and "Local Stars." This categorization motivates the title of our paper. Global giants are brands that are sold in many countries. Local stars are brands

sold in few markets, but which achieve very high market share in their country of origin. Table 3 displays Diageo's most prominent global giants and selects seven examples of local stars. The brands shown in table 3 are remarkably old, originating from 47 to 261 years ago. Not one was invented by Diageo.¹⁴ Diageo has mainly expanded its brand portfolio by acquiring brands invented long ago by other firms. The same is true for the major beer brand owners.

Global Giants								
	Johnnie Walker.	Gordon's	o Tanqueray	SMIRNOFF	Captain Morgan	BALLEYS	GUINNESS	
Origin:	UK	UK	UK	Russia	Jamaica	Ireland	Ireland	
# Markets:	68	21	28	64	43	57	30	
rank (world):	2nd	30th	46th	1st	12th	24th	21st	
born (bought):	1860 (1997)	1769 (1997)	1830 (1997)	1864 (1987)	1944 (2001)	1973 (n/a)	1759 (1997)	
	Local Stars							
	Ypi óca	McDowells		CAQUE	Bundaherg	Coun Royal .	N DA K R Fyger	
Origin:	Brazil	India	Turkey	Venezuela	Australia	Canada	Kenya	
# Markets:	2	2	2	4	1	3	1	
rank (origin):	6/44	1/47	1/51	2/44	5/119	5/87	1/14	
born (bought):	1846 (2012)	1963 (2012)	1944 (2011)	1961 (2001)	1888 (2000)	1939 (2001)	1923 (2000)	

Table 3: Diageo's Global Giants and Local Stars

Notes: Rank of Global Giants is out of 1681 spirits brands (first 6 columns) and 1567 beer brands (7th column). Rank of Local Stars shown relative to number of brands offered in the origin country. The year in () refers to acquisition by Diageo or its predecessor Grand Metropolitan.

Type of Brand:	30+ markets		Single market		#1 brand in its market		
Product	% count	% value	% count	% value	% home	# dest.*	% share*
Beer	0.3	9.7	86.9	47.0	77.6	1	24.5
Spirits	0.9	15.6	81.8	51.3	50.7	3	13.3
Carbonates	1.2	64.5	84.4	14.4	5.6	90	32.8

*: Median number of destinations and market shares of top brand.

Table 4 provides statistics on the importance of global giant and local star brands in beer and spirits (our focus) and carbonates (as a comparison). It shows that there are very few brands that sell significant amounts in 30 or more markets. While rare, global giants

¹⁴Bailey's Irish Cream was invented in 1973 within a division of Grand Metropolitan.

account for a disproportionate amount of sales. For beer and spirits, the global giants account for 10% and 16% of world sales. Soft drink giants are much more dominant, delivering 64% of world sales. Single-market brands, which constitute over 80% of brands for all three goods, are relatively unimportant in carbonates (14% of world sales) whereas they account for about half the sales of beer and spirits.¹⁵ While most single-market brands have low market shares, local stars are the leading brands in most markets. For beer, 78% of the market leaders have domestic origins (although 72% of them were foreign-owned by 2018). The lead brand's median number of destinations is just one. Their median share of the market is one quarter. This contrasts sharply with carbonates, where foreign global giants usually are the top brands. Spirits resemble beer, but the dominance of local stars is less extreme.

The salient feature of beer and spirits markets around the world is the coexistence of global giant brands with market-dominating local brands. When the owners of the former buy the latter, they have an incentive to raise markups. We now turn to the model we use to quantify how brand ownership patterns affect equilibrium markups.

3 The nested CES multi-product oligopoly model

The data described above guide the assumptions of the model. A finite number of firms compete oligopolistically, selling one or more brands in multiple markets. In addition to the firms whose market shares are listed individually (the oligopolists), our data contains an entry for a residual set of sales by small brands. As the market shares of these brands are individually less than 0.1%, we model them collectively as a monopolistically competitive fringe with exogenous mass.¹⁶ The next two subsections show how the oligopoly markups are determined, which then informs the way we obtain key elasticity parameters and back out the core concept of "brand type."

3.1 Demand

Consumers' preferences over product categories (here beer and spirits) exhibit a Constant Elasticity of Substitution (CES) η . Within these product groups, there is a lower nest of substitution between modules (e.g. lager within beer, gin within spirits) with a CES of ρ . A final CES nest considers brands within modules (e.g. Bud Light and Miller Lite),

¹⁵Appendix figure B.1 visualizes these extensive margin patterns for beer, spirits and carbonates brands.

¹⁶The mass of fringe brands can expand exogenously over time (for example, to reflect the growth of craft beers). Moreover, the sales volume of the fringe responds to markup changes by the oligopolists. Our counterfactuals do not incorporate entry/exit by the fringe in response to mergers.

which substitute with CES σ .¹⁷ This nested CES setup is one of the preference structures that Björnerstedt and Verboven (2016) used to analyze the effects of a large merger in the Swedish analgesics market.¹⁸

While the IO literature mainly uses random coefficient logit demand, the nested CES has advantages of high tractability and low data requirements that are essential for the exercises conducted in this paper. These features permit us to replicate the analysis across 76 national markets. The CES model imposes stronger restrictions on substitution elasticities than the random coefficients methods preferred in a large part of the IO literature. However, Head and Mayer (2021) show that a CES model (calibrated to replicate the observed average elasticity of substitution between brands) can do a good job of approximating aggregate outcomes of rich substitution models in counterfactual simulations.

Formally, consumers allocate a fixed beverage budget, X_n , across product groups, indexed g, with utility

$$U_n = \left[\sum_g \left(Q_{gn}^G\right)^{\frac{\eta-1}{\eta}}\right]^{\frac{\eta}{\eta-1}},\tag{1}$$

which gives the equilibrium expenditure on sector g in market n as

$$X_{gn}^{G} = \left(P_{gn}^{G}/P_{n}\right)^{1-\eta} X_{n} \quad \text{with} \quad P_{n} = \left[\sum_{g} \left(P_{gn}^{G}\right)^{1-\eta}\right]^{\frac{1}{1-\eta}},$$
(2)

where P_{gn}^G and P_n are group and overall price indices (defined below).¹⁹

Each group bundle Q_{gn}^G is itself an aggregation from modules (denoted by *m*), and the module-level quantity index Q_{mn}^M sums over brands (*b*), such that²⁰

$$Q_{gn}^{G} = \left[\sum_{m} \left(Q_{mn}^{M}\right)^{\frac{\rho-1}{\rho}}\right]^{\frac{\rho}{\rho-1}} \quad \text{and} \quad Q_{mn}^{M} = \left[\sum_{b \in \mathcal{E}_{n}} \left(A_{bn}q_{bn}\right)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}, \tag{3}$$

where q_{bn} denotes the quantity consumed of each brand b in market n, for the set of brands

¹⁷Adding a nest of substitution between products owned by the same firm would not alter the oligopoly markups (Hottman et al., 2016; Nocke and Schutz, 2018).

¹⁸Atkeson and Burstein (2008), Gaubert and Itskhoki (2018), and Burstein et al. (2019) are examples of recent work by trade economists also using Nested CES. However those papers include only two layers of CES nesting and do not consider multiproduct firms.

¹⁹The g in our application are beer, spirits and other beverages (wine, coffee, tea, soft drinks, bottled water, juice).

²⁰Each brand is implicitly associated with a unique sector g, so we dispense with g subscripts on all variables with b subscripts. Note also that while our empirics will feature several years of data, we defer the use of notation t until the point where it is indispensable.

that decided to enter market n, \mathcal{E}_n . As shown in tables 1 and 4, there is considerable cross-sectional variation in the extensive margin of where brands are offered. However, over the 12-year period of our data, there is relatively little time-series variation in \mathcal{E}_n : Appendix section B documents very low rates of adding and dropping brands across markets for beer and spirits. More crucially for our merger counterfactuals, ownership changes mainly leave intact the current patterns of where brands are offered. We corroborate this with detailed examinations of four prominent mergers in the same appendix. Since brand entry and exit do not appear to be an important aspect of the data and would prevent us from using exact hat algebra for the counterfactuals, the model treats \mathcal{E}_n as exogenous.

The market-dependent demand shifter A_{bn} , called "appeal" by Hottman et al. (2016), allows the model to capture the feature that a brand can be popular in one country (usually its origin), but be less attractive to consumers in other countries. The most of obvious interpretation of A_{bn} is simply a measure of perceived quality. In the beer and spirits industries, it seems likely that current perceived quality is the outcome of advertising campaigns. Another important process shaping A_{bn} is the accumulation of a customer base established through word of mouth and imitation. The gradual growth of loyal customers for a brand in a given market would be reflected in larger A_{bn} .²¹ Therefore our model is consistent with an advertising game or customer accumulation process in the background that determines the A_{bn} in each period. As in Sutton (1991, pp. 48–60), we compute a static Nash (Cournot and Bertrand) markup based on the perceived qualities that are taken as given in the second stage.²²

The market share of brand b in module m of market n is

$$s_{b|mn} = (p_{bn}/A_{bn})^{1-\sigma} \left(P_{mn}^{M}\right)^{\sigma-1},$$
(4)

for $b \in \mathcal{E}_n$ and 0 otherwise. We use p_{bn} to denote the price of brand *b* in market *n*. The three price indices are given by

$$P_{mn}^{M} = \left[\sum_{b \in \mathcal{E}_{n}} \left(\frac{p_{bn}}{A_{bn}}\right)^{1-\sigma}\right]^{\frac{1}{1-\sigma}}, \quad P_{gn}^{G} = \left[\sum_{m} \left(P_{mn}^{M}\right)^{1-\rho}\right]^{\frac{1}{1-\rho}}, \text{and} \quad P_{n} = \left[1 + \sum_{g} \left(P_{gn}^{G}\right)^{1-\eta}\right]^{\frac{1}{1-\eta}}.$$
 (5)

²¹Papers that have explored the consequence of customer base building include Dinlersoz and Yorukoglu (2012) and Gourio and Rudanko (2014).

²²Our interpretation leaves out two potentially important effects that we leave for future research: 1) the use of discounts to build customer bias will initially lower markups below the Static Nash level, 2) if the loyal customers have high switching costs, it will tend to make demand less price elastic.

The market share of brand b in the g sector of country n is

$$s_{b|gn} = s_{b|mn} \times s_{m|gn} = s_{b|mn} \left(P_{mn}^M / P_{gn}^G \right)^{1-\rho} = (p_{bn} / A_{bn})^{1-\sigma} (P_{mn}^M)^{\sigma-\rho} (P_{gn}^G)^{\rho-1}, \tag{6}$$

and total quantity sold by brand b in country n is

$$q_{bn} = p_{bn}^{-\sigma} A_{bn}^{\sigma-1} (P_{mn}^M)^{\sigma-\rho} (P_{gn}^G)^{\rho-\eta} (P_n)^{\eta-1} X_n.$$
(7)

Let **E** be the $N \times N$ elasticity matrix with element $\mathbf{E}_{ij} = -\frac{\partial q_i}{\partial p_j} \frac{p_j}{q_i}$. We now follow Björnerstedt and Verboven (2016) in determining the elements of this matrix, E_{ij} , where j is the brand whose price changes and is i the brand whose demand responds. Letting \mathbb{I}_x denote an indicator for when the subscripted condition x is true, the ith row and jth column of **E** are given by²³

$$E_{ij} = -\frac{\partial q_i}{\partial p_j} \frac{p_j}{q_i} = \sigma \mathbb{1}_{i=j} - (\sigma - \rho) s_{j|m} \mathbb{1}_{m(i)=m(j)} - (\rho - \eta) s_{j|g} \mathbb{1}_{g(i)=g(j)} - (\eta - 1) s_j.$$
(8)

Thus, the three indicators in this equation correspond to brands *i* and *j* being the same, belonging to same module, and belonging to the same group. For the nesting to be relevant, the lower-level substitution should be more elastic than the upper-level one: Thus $\sigma \ge \rho$ and $\rho \ge \eta \ge 1$. As a consequence, the maximum own price elasticity is σ , with rising market shares leading to lower elasticities.

3.2 Markups for different conduct assumptions

The firm maximizes the sum of brand-level profits in its portfolio (\mathcal{F}_f) of brands. Using equation (7), we obtain:

$$\Pi_{fn} = \sum_{b \in \mathcal{F}_f} q_{bn}(p_{bn} - c_{bn}) = (P_{gn}^G)^{\rho - \eta} P_n^{\eta - 1} X_n \sum_{b \in \mathcal{F}_f} (p_{bn} - c_{bn}) p_{bn}^{-\sigma} A_{bn}^{\sigma - 1} (P_{m(b)n}^M)^{\sigma - \rho}.$$
 (9)

We have added the m(b) notation in the summation to make it clear that each brand faces a different price index (the one relevant for its module). In their first order conditions, firms take X_n as exogenous but they internalize the cross-brand effects of prices that work through P^M , the effect on the group price index, P^G , and how the overall price index of beverages P_n is affected.

To express the optimal markup, it proves useful to define a cross-brand ownership

²³Aside from the new parameterization, this equation is the same as the un-numbered equation on p. 160 of Björnerstedt and Verboven (2016).

matrix. Let Θ be the brand-to-brand common ownership matrix. Using o() to denote the mapping between brand b and its owner firm f,

$$\mathbf{\Theta}_{ij} = egin{cases} 1 & ext{if } o(i) = o(j) \\ 0 & ext{otherwise.} \end{cases}$$

With this notation established we solve for the static Nash optimal markups with first price (Bertrand) and then quantity (Cournot) as the firms' decision variables.

3.2.1 Bertrand-Nash case

The first order condition for firm f, when setting prices for brand i is

$$q_{in} + \sum_{j} \Theta_{ij} \frac{\partial q_{jn}}{\partial p_{in}} (p_{jn} - c_{jn}) = 0.$$
(10)

Letting Δ be the matrix of demand derivatives, $\frac{\partial q_j}{\partial p_i}$, and suppressing the *n* subscript to focus on a single market, the additive markup form is

$$\mathbf{p} - \mathbf{c} = -(\boldsymbol{\Theta} \odot \boldsymbol{\Delta})^{-1} \cdot \mathbf{q}, \tag{11}$$

where \odot is the element-by-element Hadamard multiplication of two matrices with the same dimensions and \cdot is matrix multiplication.

It is intuitive and convenient to establish the relationship between the Lerner formulation of the markup and the demand *elasticities*. Returning to the FOC, (10), divide the whole LHS by q_i and then, within the summation, multiplying and dividing by p_j , we have

$$\frac{q_i}{q_i} + \sum_j \Theta_{ij} \frac{\partial q_j}{\partial p_i} \frac{p_j}{q_i} \frac{p_j - c_j}{p_j} = 1 + \sum_j \Theta_{ij} \frac{\partial q_i}{\partial p_j} \frac{p_j}{q_i} L_j = 1 - \sum_j \Theta_{ij} E_{ij} L_j = 0,$$
(12)

where $L_j \equiv (p_j - c_j)/p_j$ is the Lerner index and where the *second expression imposes symmetry in the demand derivatives*,²⁴ namely $\frac{\partial q_j}{\partial p_i} = \frac{\partial q_i}{\partial p_j}$. In matrix notation this is $\mathbf{1} - [\mathbf{\Theta} \odot \mathbf{E}] \cdot \mathbf{L} = \mathbf{0}$. Solving for the matrix of Lerner indexes we have

$$\mathbf{L} = (\boldsymbol{\Theta} \odot \mathbf{E})^{-1} \cdot \mathbf{1}. \tag{13}$$

²⁴Such symmetry is expected for demand curves resulting from utility maximization and it can be shown to be a feature of the nested CES demand curve.

A key feature of nested CES is that the demand elasticities are *quasi-independent* of price, as noted by Björnerstedt and Verboven (2016). Because the E_{ij} are functions of market shares only, we can derive intuition on the markup for the special case of a firm whose brands are all in the same module. The Lerner index for single-module firms simplifies to a function of *firm*-level market shares in the module, group, and beverages:²⁵

$$L_f = 1/E_f$$
, where $E_f = \sigma - (\sigma - \rho)S_{f|m} - (\rho - \eta)S_{f|g} - (\eta - 1)S_f$. (14)

Acquiring a firm that operates solely in the same module will raise all three market shares in (14), and therefore the new markup will be higher than for either firm before. Setting $\rho = \sigma$ and $S_f = 0$ delivers the Atkeson and Burstein (2008) two-layer, small-in-the-large model of CES oligopoly, with $L_f = 1/(\sigma - (\sigma - \eta)S_{f|g})$. With $S_f = 0$ and three-layer CES $(\rho < \sigma)$, equation (14) exhibits three interesting limit cases: (1) A firm that monopolizes its module has $L_f \rightarrow 1/\rho$ as the module's share of the category g goes to zero. (2) A firm that monopolizes an entire category has $L_f = 1/\eta$. (3) A firm with negligible market shares at all levels has the monopolistic competition Lerner index of $L_f = 1/\sigma$.

3.2.2 Cournot-Nash case

The first order condition for firm f, when setting quantities for brand i is

$$(p_i - c_i) + \sum_j \Theta_{ij} \frac{\partial p_j}{\partial q_i} q_j = 0.$$
(15)

One can use the similar manipulations as in the Bertrand case: divide the FOC equation by p_i ,

$$L_i + \sum_j \Theta_{ij} \frac{\partial p_j}{\partial q_i} \frac{q_j}{p_i} = L_i + \sum_j \Theta_{ij} \frac{\partial p_i}{\partial q_j} \frac{q_j}{p_i} = 0,$$
(16)

where the second equality assumes symmetry of the inverse demand derivatives, i.e. $\frac{\partial p_j}{\partial q_i} = \frac{\partial p_i}{\partial q_j}$. As the inverse of a symmetric matrix is also symmetric (Strang, 2005, p. 57), we still need just the assumption we made in the Bertrand case of symmetry in the demand derivatives. The inverse function theorem states that we can obtain the elasticities of inverse demand (multiplied by -1) as the corresponding elements of \mathbf{E}^{-1} . Using this we can express the FOC in matrix notation as $\mathbf{L} - [\boldsymbol{\Theta} \odot \mathbf{E}^{-1}] \cdot \mathbf{1} = \mathbf{0}$. Solving for the Lerner

²⁵In that case, results in (Hottman et al., 2016, appendix S2.2) imply that the term $\sum_{j} \Theta_{ij} E_{ij} L_{j}$ in equation (12) simplifies to $L_f \sum_{j} \Theta_{ij} E_{ij}$. Hence $L_f = 1/E_f$, where E_f is obtained by summing equation (8) over all brands owned by f.

index matrix for Cournot yields,

$$\mathbf{L} = (\boldsymbol{\Theta} \odot \mathbf{E}^{-1}) \cdot \mathbf{1}. \tag{17}$$

3.3 Estimating demand elasticities

The share of brand *b* in the beverages expenditures (X_n) of country *n* is

$$s_{bn} = s_{b|mn} s_{m|gn} s_{gn} = p_{bn}^{1-\sigma} A_{bn}^{\sigma-1} (P_{mn}^M)^{\sigma-\rho} (P_{gn}^G)^{\rho-\eta} (P_n)^{\eta-1}.$$
 (18)

The share of the outside good is $s_{0n} = (P_n)^{\eta-1}$. Taking the ratio s_{bn}/s_{0n} eliminates the economy-wide price index P_n , but price indexes for the *b*'s module and category remain:

$$\frac{s_{bn}}{s_{0n}} = (p_{bn}/A_{bn})^{1-\sigma} (P_{mn}^M)^{\sigma-\rho} (P_{gn}^G)^{\rho-\eta}.$$
(19)

Inverting the definitions of the conditional shares to obtain the module and category price indexes, we have

$$P_{mn}^{M} = s_{b|mn}^{1/(\sigma-1)}(p_{bn}/A_{bn}), \quad \text{and} \quad P_{gn}^{G} = s_{m|gn}^{1/(\rho-1)}P_{mn}^{M}.$$
(20)

Substituting the two price indices into equation (19), we obtain

$$\frac{s_{bn}}{s_{0n}} = (p_{bn}/A_{bn})^{1-\eta} \times s_{b|mn}^{\frac{\sigma-\eta}{\sigma-1}} \times s_{m|gn}^{\frac{\rho-\eta}{\rho-1}}.$$
(21)

Taking logs, adding time subscripts, and letting $(\eta - 1) \ln A_{bnt} = \alpha_b + x_{bn}\beta + \zeta_{bnt}$, yields the estimating equation

$$\ln s_{bnt} - \ln s_{0nt} = \alpha_b + x_{bn}\beta + (1 - \eta)\ln p_{bnt} + \frac{\sigma - \eta}{\sigma - 1}\ln s_{b|mnt} + \frac{\rho - \eta}{\rho - 1}\ln s_{m|gnt} + \zeta_{bnt},$$
(22)

where x_{bn} are brand attributes that differ depending on the market, α_b are all features of the brand that are common to all markets, and ζ_{bnt} is an error term. Equation (22), originates in Berry (1994) but was extended to three layers by Verboven (1996) and to the case of consumers with constant expenditure shares (CES) by Björnerstedt and Verboven (2016). The novel aspect of our equation is the parameterization in terms of the elasticities of substitution at each level of the nest and the inclusion of brand fixed effects in $\ln A_{bn}$. The coefficients on $\ln p_{bnt}$, $\ln s_{b|mnt}$, and $\ln s_{m|gnt}$, are denoted $-\alpha$, σ_1 , and σ_2 respectively in Björnerstedt and Verboven (2016). The top row of table 5 maps these coefficients back to the three elasticities of substitution, σ (between brands), ρ (between modules), and η (between groups).

Estimation of equation (22) requires instruments for $\ln p_{bn}$ and the two market shares. The instruments used by Björnerstedt and Verboven (2016) are counts of brands and firms by period, category and module (group and sub-group in their terminology). Our setting has the additional advantage of having many markets. This gives a spatial dimension in the x_{bn} attributes (Stella's brand home of Belgium is close to the UK market but far from the US) and thereby permits identification of the β vector even in our specification of (22) which includes brand-level fixed effects (α_b). The multi-country setting also gives an additional dimension of variation for the instruments.

Table 5 provides estimates of the nested CES critical elasticities using results from our estimation of equation (22). We provide two sets of results, one using only the brands which we were able to allocate to one of the modules of the Nielsen categories, and one also using the remainder of the brands and allocating them to an "Other Brands" module.

Elasticity: Parameter correspondence:	Brands (σ) 1 + $\alpha/(1 - \sigma_1)$		Categories (η) 1 + α
All obs.	4.99	3.49	2.82
	(1.13)	(0.51)	(0.32)
Known module	5.01	3.49	2.8
	(1.14)	(0.51)	(0.31)

Table 5: Estimated elasticities in three-layer nested CES demand

Note: Standard errors, with clustering at the brand *b* level, computed via the delta method since the elasticities are transformations of the coefficients. The variables included in x_{bn} are the frictions defined in the next section: home, distance and common language, specified based on both origin of the brand and headquarters of the owner. The instrumental variables are counts of brands by *nt*, *gnt*, *fnt*, *fgnt*, *fgmnt* and *gmnt*. There are fixed effects for each brand *b*.

To verify these parameter estimates are reasonable, we first compute the implied average own-price elasticities using equation (8). This average E_{ii} equals 4.81 with a standard deviation of 0.33. We can compare this to the reported mean or median brand elasticities from the industrial organization literature studying the beer and spirits industries. In all, we have collected 18 estimates: the average is 3.93, with a standard deviation of 0.96.²⁶ While the literature average estimate is 18% lower than those we obtain, our estimates derive from a much larger set of countries (76). Considering estimates from the five papers on the beer industry, the average elasticity is 4.48, which is even closer to our average

²⁶The estimates come from Asker (2016), De Loecker and Scott (2016), Hausman et al. (1994), Miller and Weinberg (2017), Pinkse and Slade (2004) for beer, and Miravete et al. (2018) and Conlon and Rao (2015) for spirits.

estimate of 4.81.

4 Estimation of ownership effects on brand performance

The focus in this section is to estimate the impact of firm ownership on brand performance (measured as market share, appeal, and cost-adjusted appeal). We consider both a pure ownership effect, i.e. the way an individual firm improves a brand's performance everywhere, and a localized effect that depends on the proximity of the firm's HQ to each market served by the brand. To isolate these two ways that the owner of a brand matters, we need to control for factors that operate at the brand level. Here again, there are two aspects: the global brand appeal and the differential appeal associated with proximity between the brand's origin and the market where it is being sold.

In the coming derivation of the estimating equations, there are three mappings that we use repeatedly for the definitions of variables and econometric specifications:

- o(b, t) maps a brands to its *owner* in year t.²⁷
- h(f) maps a firm to location of its *headquarters*.
- *i*(*b*) maps a brand to its *origin*, the country where the brand was introduced.

4.1 Backing out cost-adjusted appeal (brand type)

Borrowing from Nocke and Schutz (2018), the term "brand type" refers to the attribute that determines a brand's market share. Denote it φ following Melitz (2003) footnote 7 pointing out that firm heterogeneity could be isomorphically represented as either a demand shifter or physical productivity. In terms of determining equilibrium brand market shares, all that matters in the CES model is the ratio, $\varphi_{bn} \equiv A_{bn}/c_{bn}$, which we refer to as cost-adjusted appeal or brand type.²⁸ With estimates of the demand elasticities, data on brand sales shares in a market allow us to back out all the φ_{bnt} up to a normalization.

The true cost-adjusted appeal, φ_{bnt} , is unobserved but the equilibrium quantities and prices only depend on the value of φ_{bnt} relative to its geometric mean within the module and market, $\tilde{\varphi}_{m(b)nt}$. Our inferred brand types will thus have the same normalization and

²⁷There are some brands, e.g. Fosters, whose owner varies across countries. We omit the *n* subscript from o(b, t, n) in the notation, but take it into account in the estimation and counterfactuals.

²⁸Melitz (2003) points out the isomorphism in a model of CES single-variety monopolistic competition. Nocke and Schutz (2018) generalize it to multi-product oligopoly and also show that a similar isomorphism applies in the logit model with the φ expressed as a *difference* between appeal and cost.

notation as Hottman et al. (2016) use for inferring brand appeal. We invert the market share of brand b in module m of market n (equation 4) to obtain:

$$\check{\varphi}_{bnt} = \frac{\varphi_{bn}}{\tilde{\varphi}_{m(b)nt}} = \left(\frac{s_{b|mnt}}{\tilde{s}_{m(b)nt}}\right)^{1/(\sigma-1)} \frac{\mu_{bnt}}{\tilde{\mu}_{m(b)nt}},\tag{23}$$

where $\tilde{s}_{m(b)nt}$ is the geometric mean of the market shares of the brands within the module $(s_{b|mnt})$ for all brands active in the same module as b in a given country and year. The market shares are directly observable and the markups are implied by the underlying model of market structure (section 3.2 providing formulas for both Bertrand and Cournot cases).

In a way similar to Khandelwal et al. (2013) and Redding and Weinstein (2018), one can infer the relative demand shifter from observables in the data (market shares and prices) combined with an estimate of σ :

$$\check{A}_{bnt} = \frac{A_{bnt}}{\tilde{A}_{m(b)nt}} = \left(\frac{s_{b|mnt}}{\tilde{s}_{m(b)nt}}\right)^{1/(\sigma-1)} \frac{p_{bnt}}{\tilde{p}_{m(b)nt}}.$$
(24)

Note that unlike brand type, brand appeal can be backed out without imposing a conduct assumption. However, inferring brand appeal does require price data. For both φ_{bnt} and A_{bnt} we can only identify the parameters within a product-market-year. Intuitively, multiplying all the φ_{bnt} or A_{bnt} by a scalar would not change the equilibrium market shares conditional on the other variables.

4.2 Estimating equations

We decompose inferred brand type ($\check{\varphi}_{bn}$) as the product of five factors: the brand's costadjusted appeal common to all markets, φ_b^B , the firm's cost-adjusted appeal, $\varphi_{o(b,t)}^F$, the frictions at the level of brand origins, $\mathbf{X}_{i(b)n}$ and headquarters, $\mathbf{X}_{h(o(b,t))n}$, and an error term, ε_{bn} . The final estimating equation for cost-adjusted appeal accounts for the fact that we have multiple years of data (hence the index *t*) and uses our inferred values, $\check{\varphi}_{bnt}$, in place of the unobserved φ_{bnt} :

$$\ln \check{\varphi}_{bnt} = -\ln \tilde{\varphi}_{m(b)nt} + \ln \varphi_b^B + \ln \varphi_{o(b,t)}^F + \mathbf{X}'_{i(b)n} \mathbf{d}^B + \mathbf{X}'_{h(o(b,t))n} \mathbf{d}^F + \varepsilon_{bnt}.$$
 (25)

On the right-hand-side of equation (25), $\mathbf{X}_{i(b)n}$ and $\mathbf{X}_{h(o(b,t))n}$ capture the impact of observable frictions on φ_{bnt} . They include effects such as home bias in preferences, which enters via A_{bnt} , as well as costs of distributing remotely, which would enter via c_{bnt} . We focus on two "home" variables as determinants of $\mathbf{X}_{i(b)n}^B$ and $\mathbf{X}_{h(o(b,t))n}^F$. The first, home_{i(b)n}, takes a value of 1 for brands sold in their country of origin (i = n). The second, home_{h(o(b,t))n}, equals one when the current owner of a brand has its headquarters in the market (h = n). We also include common language and the log of distance, with *in* and *hn* formulations for each variable.

The other key determinants of $\ln \tilde{\varphi}_{bnt}$ are three sets of fixed effects. The first one captures all determinants that are related to module m in market n and time t, with theoretical correspondence $-\ln \tilde{\varphi}_{m(b)nt}$. The two other sets of fixed effects are the most important for our analysis: the brand and owner effects $(\ln \varphi_b^B \text{ and } \ln \varphi_{o(b,t)}^F)$. As is well known since Abowd et al. (1999), estimation of the effect of individual owners requires "mobility" of brands across firms. Thisalso explains why the time subscript is actually indispensable when estimating equation (25).

4.3 Baseline estimation results

Table 6 reports results for regressions that pool beer and spirits brands. The last two columns provide estimates of determinants of brand type, as specified in equation (25). Columns (1) and (2) use the same set of determinants, with dependent variables changed as being the market share (s) and appeal (A) of brand b in market-year nt.

The most striking result is the huge advantages held by home-origin brands. Since exp(1.148) = 3.15, home increases market share of a brand within its module by 215%. Our estimate of the home advantage for beer and spirits brands is larger than the 126% estimate for car brands obtained in Head and Mayer (2019). Distance from brand origin also reduces market share, with an elasticity of -0.25. Head and Mayer (2019) estimate a somewhat larger elasticity of -0.34 for cars.

The market share effects combine cost and appeal effects with the substitution elasticity. The pure effect of being a home brand on cost-adjusted appeal is equivalent to a $\exp(0.308) - 1 = 36\%$ price discount (Bertrand conduct). The majority of this comes from the taste side (home bias). In particular, being a home brand raises demand by an amount corresponding to a $\exp(0.176) - 1 = 19\%$ price reduction.²⁹

Our data allow us to mention a novel dimension of home bias: whether home ownership improves sales performance of a brand, either via higher appeal or lower costs. This effect is central when considering multinational acquisitions since they can transform a local star brand into one that is foreign-owned. Column (1) of Table 6 shows that

²⁹Goldberg and Verboven (2001) and Coşar et al. (2018) find significant home bias attributable to preferences in the car industry but functional form differences make it hard to compare their parameter estimates to ours.

in the raw data there is indeed a large increase in market share for home-owned brands although not as large as the effect of brand origin. Moving to columns (3) and (4), the estimates imply that the underlying cost of foreign ownership corresponds to an 11 to 12 percent reduction in cost-adjusted appeal.³⁰

Table 0. Drand performance regressions. Deer and opints							
			Bertrand	Cournot			
dependent variable:	$\ln s_{b mnt}$	$\ln \check{A}_{bnt}$	$\ln \hat{\varphi}$	$\hat{\phi}_{bnt}$			
1	mkt. shr	appeal	appea				
home	1.148^{a}	0.176^{a}	0.308^{a}	0.319 ^a			
	(0.132)	(0.054)	(0.035)	(0.036)			
		· · ·					
distance	-0.254^{a}	-0.031	-0.065^{a}	-0.067^{a}			
	(0.044)	(0.020)	(0.012)	(0.012)			
	```'	```	``'	``'			
common language	$0.212^{b}$	0.031	$0.055^{b}$	$0.056^{b}$			
0 0	(0.087)	(0.039)	(0.023)	(0.023)			
		· · ·					
home (HQ)	$0.360^{a}$	$0.092^{b}$	$0.115^{a}$	$0.128^{a}$			
	(0.121)	(0.045)	(0.032)	(0.033)			
	· · · ·	× ,	<b>``</b>	· · /			
distance (HQ)	0.047	0.026 ^c	0.012	0.012			
	(0.036)	(0.014)	(0.009)	(0.010)			
	~ /	× /	× ,	· · · ·			
com. lang. (HQ)	0.098	0.004	0.027	0.029			
$0 < \sim$	(0.065)	(0.027)	(0.018)	(0.018)			
	()	()	()	(			
Observations	95,299	95,299	95,299	95,299			
$R^2$	0.817	0.723	0.656	0.660			
1	0.017	0.723	0.000	0.000			

Table 6: Brand performance regressions: Beer and Spirits

Notes: Standard errors in (), clustered by origin-market dyads. Fixed effects at the firm, brand-product and market-year-product dimensions included in each specification. HQ variables defined with respect to brand owner's headquarters country. Significance levels: 1% (*a*), 5% (*b*), and 10% (*c*).

The  $R^2$  of the brand type estimations in table 6 round to 0.66 for both conduct assumptions, indicating that idiosyncratic shocks explain about 34% of the variation in  $\ln \check{\varphi}_{bnt}$ . This finding motivates the usefulness of exact hat algebra (EHA) for counterfactuals since it implicitly takes into account the unobserved determinants of market share that are invariant to the counterfactual.

The regressions in Table 6 estimate the effect of frictions averaging over 12 years, pooling beer and spirits. To assess how home biases differ across these groups, and how

³⁰The calculations are  $\exp(-0.115) - 1 = -0.109$  for Bertrand, and  $\exp(-0.128) - 1 = -0.120$  for Cournot.

they evolve over time, we estimate a model for beer and spirits separately, interacting the home origin and HQ dummies with year dummies. Figure 4 graphs the results, expressed as *ad-valorem* equivalents (AVE) of the home advantage for brand type ( $\varphi$ ).³¹ The home bias estimated under the Cournot conduct assumption is systematically larger than under Bertrand. The graph displays the range between the two estimates using blue (origin) and red (HQ) ribbons. We use the same coloring schemes (with symbol-separated lines) to display the AVEs of the part of home bias that comes from appeal. These appeal effects do not depend on conduct, since they are extracted directly as demand shifters.

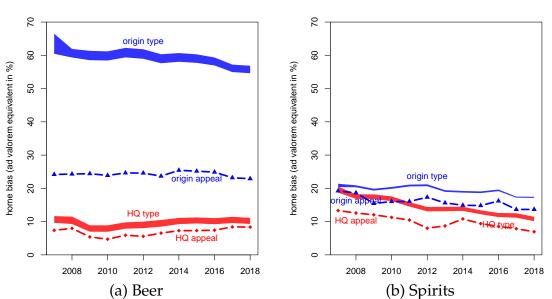


Figure 4: The evolution of different forms of home brand advantage

Note: Upper and lower bounds of each "ribbon" use Cournot and Bertrand markup assumptions, respectively.

As seen in panel (a) of Figure 4, the total effect of being a home origin beer brand is equivalent to a 55–65% tax imposed on foreign-origin competitors. This large home bias helps us understand the existence of the local stars phenomenon. Even if they lack universal appeal (which explains why they rarely sell in other markets), domestic brands can achieve very large home market shares under this estimated level of protection from foreign competition. As a consequence, foreign firms find it difficult to penetrate the market without purchasing those local stars.

For beer brewers, the consumer preference for domestic brands (a 25% AVE) accounts for about one third of the home origin type advantage. The AVE of the consumer bias is almost the same in spirits (panel b). For that product, it represents a much larger share of overall home advantage in cost-adjusted appeal. A natural explanation is that spirits

³¹The formula is  $100 \times [\exp(\mathbf{d}) - 1]$ , where **d** is the home coefficient in the brand type ( $\varphi$ ) regression.

have a much larger value-to-weight ratio. To the extent that domestic-origin brands are also produced locally, transport costs incurred by foreign brands should matter more for beer.

HQ-related home advantage is estimated as equivalent to about a 10% tariff for beer and spirits in 2018. This is the immediate cost increase or appeal decline imposed on a brand when bought by a foreign company. Our estimation can identify this effect, even controlling for home origin effects, from brands whose owner changes lead to a change in headquarters. To rationalize acquisitions that transfer headquarters abroad, there would need to be some gain to offset the estimated penalty of foreign ownership. The two candidates we consider are firm value-added to brand performance and increased market power.

To estimate the value-added of firms, we consider the firm-level fixed effects that form part of our regression specification. The difference between the seller and buyer firm fixed effects measures the change in cost-adjusted appeal of the brand (in all destinations) when changing owner. The structural interpretation of the estimated owner fixed effect in equation (25) is  $\ln \varphi_{o(b,t)}^F$ . A transfer of *b* to a new owner in period t + 1, raises cost-adjusted appeal by  $\ln \varphi_{o(b,t+1)}^F - \ln \varphi_{o(b,t)}^F$ . Substantial variance in the estimated firm-level fixed effects is a necessary condition for firms to add value. However, it is not a sufficient condition. In addition, brands should move from poor to strong firms. In the next subsection, we measure the variance of firm fixed effects and depict the distribution of changes in fixed effects brought about by ownership changes.

## 4.4 Estimating the contribution of firm effects

Before assessing the relative contribution of brand and firm fixed effects, we need to establish how these parameters can be separately identified. As is the case with firm and worker effects on wages, identification requires "mobility." In our context, movements are changes in the ownership of brands which connect different firms. This is analogous to how workers changing jobs connect establishments in the seminal paper by Abowd et al. (1999), now known by the initials AKM. Another helpful analogy is the literature on the value-added of teachers. As with brand owners, that literature can estimate fixed effects only for sets of teachers who are connected by in-common students.

The labor economics research estimating wage equations using worker and firm fixed effects has grappled with two main issues relating to worker mobility, both of which apply in our setting. The first concern is endogenous assignment of workers to firms. The second concern is even if mobility is conditionally exogenous, there may be too little mobility in practice to obtain unbiased estimates of the variance and correlations of the fixed effects. We now consider how these issues apply to our setting.

#### 4.4.1 Endogenous mobility bias

Adapting equation (2) in Bonhomme et al. (2023), we express the orthogonality condition needed for equation (25) to be correctly specified as

$$\mathbb{E}[\varepsilon_{bnt}|\mathbf{X}_{i(b)n}, \mathbf{X}_{h(o(b,t))n}, \varphi_b^B, \varphi_f^F, \tilde{\varphi}_{m(b)nt}, \mathbf{\Omega}] = 0,$$
(26)

where  $\Omega$  is an array defining the owner of each brand in each period. The entries for each element of this array are given by

$$\Omega_{bft} = \begin{cases} 1 & \text{if } o(b,t) = f \\ 0 & \text{otherwise.} \end{cases}$$

Adapting equation (5) of Card et al. (2013), a sufficient condition for the assignment process to ensure that  $\varepsilon$  is conditionally exogenous is given by

$$\operatorname{Prob}(o(b,t) = f|\varepsilon) = \operatorname{Prob}(o(b,t) = f) = G_{ft}(\varphi_b^B, \varphi_f^F, \tilde{\varphi}_{m(b)nt}), \quad \forall \ b,t$$
(27)

This condition allows for the assignment of brands to firms to depend very generally on the brand and firm components of cost-adjusted appeal,  $\varphi_b^B$  and  $\varphi^F$ . For example good brands may be more likely to be acquired by good firms. Also as implied by the subscripts of the function  $G_{ft}()$ , equation (27) allows some firms with better access to capital markets to be more likely to acquire brands.

There are possible violations of equation (27) that would have the potential to bias our estimates.³² First among these would be to incorporate an idiosyncratic brand-firm component into  $\varepsilon$ . If this term, denoted  $\xi_{bf}$ , has zero mean conditional on the covariates, as in Card et al. (2013), then the AKM specification in equation (25) remains valid. However, if firms observe  $\xi_{bf}$  before assignment of brands, then equation (27) would be violated. Thus, it would pose a problem if acquired brands are particularly good fits for the acquiring company or a poor fit for the selling company.

We address the endogenous mobility concern in two ways suggested by Card et al. (2013). First, we add a new specification that replaces the additive brand b and firm f fixed effects with interactive brand-firm (bf) fixed effects. If firm-brand "match effects"

³²Appendix C provides a simulation-based quantification of the bias generated by endogenous mobility in our implementation of this type of regression.

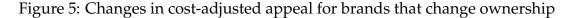
are important in determining which firms own which brands, there is a potential for bias because the error term in the additive specification could be correlated with the friction determinant or firm fixed effects. While Card et al. (2013) have only time-series wage variation to identify the worker-firm interactions, our context has the benefit of crossmarket and cross-time variation to estimate the brand-firm effects. Results displayed in columns (3) and (6) of Table 7 show that the additional fit from these interactive effects is small in our context, when compared to columns (1) and (4) showing our baseline regressions. The table also shows that the inclusion of the *bf* effects has negligible effects on the friction estimates, suggesting that the orthogonality assumption for the match effects is not strongly violated.

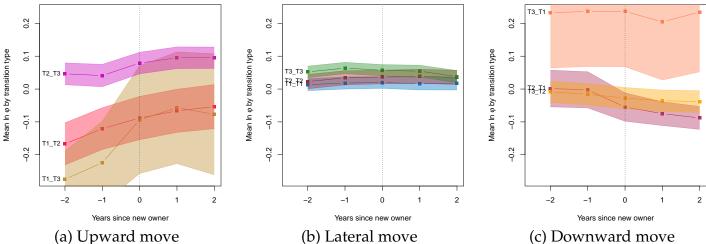
Table 7. brand type regressions with alternative neterogeneity assumptions							
		Beer		Spirits			
	(1)	(2)	(3)	(4)	(5)	(6)	
Fixed effects:	b+f	b + k(f)	bf	b+f	b + k(f)	bf	
home	$0.457^{a}$	$0.466^{a}$	$0.466^{a}$	$0.176^{a}$	$0.170^{a}$	$0.176^{a}$	
	(0.050)	(0.048)	(0.050)	(0.043)	(0.042)	(0.044)	
distance	$-0.049^{a}$	$-0.052^{a}$	$-0.055^{a}$	$-0.065^{a}$	$-0.062^{a}$	$-0.065^{a}$	
	(0.018)	(0.016)	(0.018)	(0.015)	(0.015)	(0.015)	
common language	$0.099^{b}$	$0.101^{b}$	$0.096^{b}$	0.031	0.035	0.034	
	(0.041)	(0.040)	(0.041)	(0.026)	(0.025)	(0.026)	
home (HQ)	0.082	0.054	0.072	$0.131^{a}$	$0.120^{a}$	$0.143^{a}$	
	(0.054)	(0.043)	(0.058)	(0.038)	(0.035)	(0.041)	
distance (HQ)	$-0.034^{b}$	$-0.025^{b}$	$-0.035^{c}$	$0.031^{a}$	$0.026^{a}$	$0.035^{a}$	
	(0.015)	(0.010)	(0.018)	(0.011)	(0.010)	(0.012)	
com. lang. (HQ)	-0.036	-0.035	-0.031	$0.050^{b}$	$0.042^{b}$	$0.050^{b}$	
	(0.034)	(0.031)	(0.037)	(0.020)	(0.019)	(0.020)	
Observations	34,675	34,675	34,675	60,624	60,624	60,624	
$\mathbb{R}^2$	0.733	0.729	0.748	0.603	0.596	0.608	
RMSE	0.203	0.203	0.198	0.208	0.208	0.206	

Table 7: Brand type regressions with alternative heterogeneity assumptions

Standard errors in (), clustered by origin-market dyads. Dependent variable is  $\ln \tilde{\varphi}_{bn}$  based on Bertrand conduct (Results assuming Cournot conduct shown in Table E.6 ). Market-yearproduct fixed effects in each regression. HQ variables determined by brand owner's headquarters country. In the second and fifth columns, *k* corresponds to the group FE (*K* = 10). Significance levels: 1% (*a*), 5% (*b*), and 10% (*c*).

The second approach to assess the importance of endogenous mobility is to construct an event-study analogous to the one carried out in figure VII of Card et al. (2013). The authors argue that when assignment depends on the idiosyncratic match, then the expected change in wages for movers is positive. In our context that would imply a positive





Note: T1, T2, and T3 are the terciles of the firms in the largest connected sets for beer and spirits. Each filled interval shows the mean plus or minus one standard error. Equal vertical ranges for each panel correspond to the range in means for all transitioners.

expected change in  $\varphi_{bnt}$  for brands that change owners even if the fixed effects ( $\varphi_f^F$ ) of the owners are the same. The middle panel (b) of figure 5 shows that, contrary to the prediction of idiosyncratic sorting, brands do not on average improve their  $\varphi_{bnt}$  when moving between firms in the same tercile. We do see rises for "upgrading" moves (T1_T2, T2_T3) in the left panel (a) and average declines for brands moving from the second (T2) to the lowest tercile (T3) in the right panel (c). These changes are expected with the assignment process of equation (27) which is sufficient for the validity of the estimating equation. Unlike the employer-employee datasets, our data lacks large counts of movers for every quantile transition pair. In particular, the T3 to T1 and T1 to T3 movement types are estimated based on fewer than 10 brands and are thus too noisy to infer clear patterns.

### 4.4.2 Limited mobility bias

The employer-employee and teacher-student literatures have highlighted several important lessons that are applicable to our estimation of brand and owner effects. First, the presence of firm fixed effects should not bias the estimation of the friction coefficients (home, distance, language) in Table 6.³³ Second, firm fixed effects are estimated relative to a reference firm, with *a different reference firm for each connected set*. It is therefore meaningless to compare firm fixed effects across sets or to estimate the overall variance of fixed effects. The third point coming from the AKM literature is that even within the connected

³³The coefficients are similar (differing mainly in the second decimal, and by less than a standard error) to those reported in Table E.1, which is estimated without firm fixed effects.

set, the fixed effects are often noisily measured. The reason for this has come to be termed "limited mobility bias." When few workers connect firms, Andrews et al. (2008) find that the variance of the fixed effects will be over-estimated and spurious negative correlations can appear between worker and employer fixed effects.

Jochmans and Weidner (2019) recast the concern over limited mobility as a network problem. Starting from a bipartite network—teachers and students in their example— one constructs the induced teacher-to-teacher network weighting the edges by the number of student-course combinations shared by each teacher pair (the edges in the induced graph). They show that the amount of excess variance in the teacher fixed effect estimates will be bounded from above by a function of a particular measure of the global connectivity of the induced network. This measure, denoted  $\lambda_2$ , is calculated as the smallest non-zero eigenvalue of the normalized weighted Laplacian of the induced network.³⁴ In our context, a firm whose brands have never been owned by any other firm is disconnected from other firms. Brands with multiple owners, in time or space, connect firms. But it may be that the network is only barely connected, i.e. loss of a few brands would break it into disjoint components. Figure D.1 (a) and (b) in the appendix illustrate this possibility using a graph featuring 12 firms and 12 brands. In that example, a single brand (Fosters) is critical for maintaining the connection between two sub-graphs.

When graphs are poorly connected, AKM estimates of fixed effects exhibit excess variance. This is important for us because it implies that naive AKM estimation overstates the value firms add to brands. We therefore apply three methods to mitigate this problem. The first method comes from Andrews et al. (2008), hereafter AGSU. They show that in labor data one can avoid excess variance and spurious negative correlations between worker and plant fixed effects by restricting the sample to movers (workers who change plants) and "high mobility" plants. In their context, high mobility is achieved by plants with 30 or more moving workers. AGSU assign the workers at low-mobility plants to a single "superplant" fixed effect. In our case, movers are brands who change ownership and high mobility refers to firms with ten or more brands that change ownership. Brands owned by low-mobility firms receive the same "superfirm" fixed effect.

The second method for mitigating limited mobility bias comes from Bonhomme et al. (2019), hereafter BLM. While the focus of their paper is a random effects specification, the authors report that a *group fixed effects* specification achieves similar reductions of the bias in the variance of fixed effects. The first step of this method is to group firms using k-means clustering, based on the distribution of market shares achieved by the

³⁴Appendix section D provides greater detail on this procedure.

brands the firm owns in the first period (2007 for most firms).³⁵ The second step is to re-estimate equation (25) with firm-group k fixed effect replacing firm f fixed effects. The results are shown in table 7, columns (2) and (5). Although group effects work by reducing dimensionality, using group k fixed effects instead of firm f effects lowers  $R^2$  by just 0.004 for beer and 0.007 for spirits. The k fixed-effect specification also obtains very similar estimates of the frictions. This specification provides the friction coefficients and firm group effects we use in the counterfactuals conducted in the next section.

In both of the above methods, the fundamental idea is to estimate fewer fixed effects so as to ensure that those fixed effects are for well-connected entities. Kline et al. (2020), hereafter KSS, offer a third way to estimate the variance share of fixed effects that does not restrict the dimensionality to clusters as in BLM. Instead, the KSS method consistently estimates the variance components for the original high-dimensional entities. The first step of KSS reduces the set of firms to those who remain connected to each other no matter which brand is removed. Using KSS terminology, there are no "bottleneck" brands in this restricted sample. This leave-out-match data set has 50 firms in the beer category and 43 in spirits.³⁶ The second step of KSS constructs a finite sample unbiased variance estimator that is computed by repeatedly leaving out a single match between a brand and firm.

Table 8 summarizes our results on the firm effects for beer and spirits.³⁷ Column (1) shows the incremental  $R^2$  for firm fixed effects in the full sample is just 0.005 for Beer and 0.008 for Spirits. That is, firms add very little explanatory power to a specification that already includes brand effects and the six friction variables. This is true across all the specifications considered in this table.

The standard way to measure firm value added is shown in column (4): the variance of the firm fixed effects divided by the variance of the outcome variable ( $\ln \phi_{bnt}$  here). In the full sample, this variance is meaningless (hence the NA) because there are multiple references. The largest connected set (LCS), i.e. the one in which all firm fixed effects are estimated relative to the same firm, is reported in the second row for each group. The LCS firms account for the majority of world sales.³⁸ The variance share for firms in the LCS is 12% for beer and 37% spirits. These figures overstate the importance of firms and misleadingly imply strongly negative assortative matching between owners and brands. The reason that LCS results cannot be trusted is that the  $\lambda_2$  connectivity of both sets is just

³⁵As in BLM, the features used in the clustering of firms are binned percentiles. Whereas they used 20 bins of the log wage distribution, we use five bins of  $\ln s_{bn}$ . Our use of fewer bins reflects the smaller number of brand-market observations per firm (about 6) than worker observations per establishment (about 37).

³⁶In the network illustrated in figure D.1, Fosters is a bottleneck brand.

³⁷Results for Cournot conduct are very similar to the Bertrand results shown here, so they are relegated to Appendix table E.5.

³⁸Table D.1 shows that firms in the LCS accounts for 80% of beer sales and 58% of spirits.

	(1)	(2)	(3)	(4)	(5)
Type of FE	# of FE	$\lambda_2$	$\Delta R^2$	Varshr	FE Corr
	Beer				
Firms (All)	464	0.000	0.005	NA	NA
Firms (Largest connected set, AKM)	90	0.013	0.005	0.118	-0.308
Firms (Leave-out-match, KSS)	50	0.072	0.004	0.096	-0.293
Firms (High mobility, AGSU)	22	0.171	0.003	0.041	-0.110
Clusters (BLM)	15	0.548	0.001	0.007	0.112
Clusters (BLM)	10	0.748	0.001	0.007	0.092
Clusters (BLM)	5	0.957	0.000	0.002	0.121
S	Spirits				
Firms (All)	849	0.000	0.008	NA	NA
Firms (Largest connected set, AKM)	93	0.013	0.008	0.365	-0.450
Firms (Leave-out-match, KSS)	43	0.015	0.003	0.225	-0.356
Firms (High mobility, AGSU)	19	0.071	0.006	0.096	-0.252
Clusters (BLM)	15	0.348	0.001	0.023	0.148
Clusters (BLM)	10	0.601	0.001	0.013	0.094
Clusters (BLM)	5	0.869	0.000	0.005	0.268

Table 8: The explanatory power of owner fixed effects

Notes: # of FE is either number of firms or clusters.  $\lambda_2$  measures network connectivity.  $\Delta R^2$  is the difference in  $R^2$  between the full specification and one excluding firm/cluster fixed effects. Varshr is the ratio of the variance of firm/cluster FEs to the variance of brand type (ln  $\phi_{bn}$ , conduct =Bertrand). FE corr is the correlation between brand and firm/cluster FEs. References for AKM, AGSU, BLM, KSS given in text.

0.01, compared to  $\lambda_2 = 1.00$  for a fully connected network.³⁹ Our first attempt to mitigate limited mobility bias is the KSS leave-out-match estimator. Using this method reduces the variance share of firm fixed effects, and moves the fixed effect correlation towards zero.⁴⁰ However, the degree of connectivity ( $\lambda_2 = 0.015$ ) remains too low to trust the spirits results.

The subsequent rows of Table 8 establish that when  $\lambda_2 > 0.1$ , the variance share of firm fixed effects shrinks to less than 5% for both groups. Restricting to the set of high mobility firms raises  $\lambda_2$  to 0.17 for beer and 0.07 for spirits, which is sufficient to put the variance share at 4% and 10%, for beer and spirits, respectively.⁴¹ The firm effects estimated in this sample have much smaller negative correlations with their corresponding brand effects. The variance shares shown in the Clusters (BLM) rows of Table 8 convey a common message about firm value-added whether we use K = 10 as in BLM, K = 5, or K = 15. For each K and for both beer and spirits, connectivity is over 0.3 and the variance share for firms is under 2.3%. The group fixed effects method also eliminates the negative FE correlations that appeared to point negative assortative matching.

#### 4.4.3 The impact of ownership change on cost-adjusted appeal

Figure 6 visualizes the distributions of changes in  $\check{\varphi}_{bn}$  that our estimates imply to have occurred as a consequence of the observed set of brand ownership changes. The blue densities shows changes in  $\check{\varphi}_{bn}$  attributable to changes in the headquarter country after cross-border acquisitions take place. Since there are many same-country mergers, there is an important mode at zero. The second mode (at around -0.15) corresponds to domestic brands being acquired by foreign firms. The reverse phenomena—an increase in costadjusted appeal when domestic firms purchase foreign-owned brands—is rare.

The red densities in Figure 6 show the effect of changing owners for firms in the largest connected set (LCS). The red density in the lower row of graphs is for firm-clusters (BLM, K = 10). The density has a strong peak near zero in every case, but it is especially high density for the firm-cluster fixed effects. Under group effects, the new owner frequently comes from the same group as the original one. For example, AB InBev was in the same group as SAB Miller. The difference in  $\ln \tilde{\varphi}_{bn}$  between the groups to which AB InBev and Anheuser Busch respectively belong corresponds to a small 0.04 log point improvement

³⁹Interestingly, the firm-to-firm network here is slightly more connected than the  $\lambda_2 = 0.004$  in the teacher-to-teacher network examined by Jochmans and Weidner (2019).

⁴⁰We implement the KSS estimator using the Bonhomme et al. (2023) python package **pytwoway** using the leave-out-match option.

⁴¹The firms remaining in the high mobility set still account for a respectable 71% of beer sales and 42% of spirits sales.

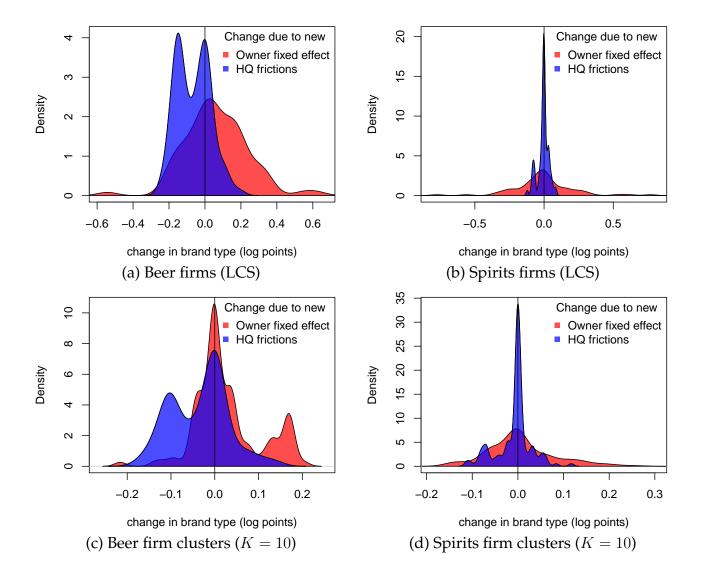


Figure 6: How ownership changes affect brand type ( $\varphi_{bn}$ )

of Budweiser's brand type. On the other hand, when AB Inbev bought London-based Camden Town Brewery in 2015, the latter benefited from a 0.14 improvement in  $\ln \check{\varphi}_{bn}$ . The Belgian craft brewery Bosteels made an even larger move (0.17) when AB InBev acquired it in 2016.⁴² Another important finding displayed in figure 6 is that the range of group effects is about 0.4 for firm-clusters in beer which is much smaller than the 1.2 range for firm effects, just as predicted by low mobility bias. A similar range shrinkage occurs for spirits.

Our results echo the findings of Blonigen and Pierce (2016), who find little evidence that mergers affect plant-level productivity. They are also in line with the Kwoka (2014) survey of 41 different mergers where only one in four cases exhibited clear performance improvements following a merger. More recently, Ashenfelter et al. (2015) and Miller and Weinberg (2017), estimate that shipping cost savings from the MillerCoors joint venture lower US prices by 2% (offsetting the price increase induced by higher concentration).

There is an important consequence of our regressions in interpreting the role of firms in the beer and spirits industries. Since firm effects contribute so little to brand performance, we see little evidence of significant marginal cost or appeal synergies in the brand amalgamation process. This raises the question of why firms find it profitable to collect brands. The obvious explanation coming from recent critiques emphasizing rising market power, and formalized within our model, is that mergers suppress competition between brands. An additional explanation would be synergies that take the form of fixed costs reductions. We quantify the market power effects on profits and determine when fixed cost reductions are needed to rationalize mergers at the end of the counterfactual exercises conducted in the next section.

## 5 Counterfactual merger policies and consumer welfare

Mergers and acquisitions of beer and spirits makers have expanded the sets of brands under the ownership of the largest multinationals (as seen in figures 1 and 2). To quantify the consequences for consumer welfare of multinational brand amalgamation, we consider counterfactual ownership configurations. Our first set of counterfactuals investigates the consumer surplus saved by antitrust remedies and foregone in less interventionist countries. We then calculate the changes in concentration and consumer surplus implied by a counterfactual scenario banning all acquisitions from 2007 to 2018.

⁴²The Bosteels-owned brand in GMID, Triple Karmeliet, won the World Beer Awards in 2008 so it seems likely the rise in  $\varphi$  came from more efficient production processes or more intensive advertising as opposed to a pure change in quality.

In addition to taking into account how alternative ownership patterns affect firm level market shares and hence their optimal markups, we also account for the changes in brand type ( $\varphi_{bn}$ ) implied by the counterfactual ownership, using friction estimates displayed in columns 2 and 5 (beer and spirits, respectively) of Tables 7 (Bertrand) or E.6 (Cournot), and firm-group k fixed effects, illustrated in Figure 6(c) and (d). The results include the difference in the estimated group fixed effect corresponding to the actual and counterfactual owners. The simulations also include the changes in frictions that are estimated to result from any ownership change that moves headquarters out of the country in question, further away, or to a country with a different language. The next subsection describes the method used for all the counterfactual computations.

### 5.1 Exact Hat Algebra (EHA) for M&A

The counterfactual stipulates a new set of brand portfolios for each firm. Firm market shares adjust to new ownership sets and to changes in brand market shares entailed by rearranging ownership, altering first-order conditions for pricing. Those first-order conditions (equations 13 for Bertrand and 17 for Cournot) depend crucially on the aggregation of brand-level market shares through demand elasticity (8). Therefore, ownership changes that we simulate imply a change in the equilibrium firm-destination market shares for three reasons: 1) the changes in the number and identity of brands owned, 2) the changes in  $\varphi_{bn}$  following ownership changes of brands, 3) the changes in equilibrium market shares of those brands resulting from 1) and 2).

The key element of EHA is to show that, in a counterfactual, one can compute the new value of a market share using the initial observed market share to capture all the unobservables that will remain unaffected by the experiment. This property has been largely exploited in the trade literature using CES demand with constant markups (Costinot and Rodriguez-Clare, 2014). We now show that this property remains with 3-level CES demand combined with oligopoly. So far as we know, this is the first application extending EHA to incorporate oligopoly markup adjustment, which permits counterfactual merger analysis. In terms of notation, we follow the established convention, and denote counterfactual equilibrium of variable x with x', and the proportional change as  $\hat{x} = x'/x$ . The counterfactual calculation procedure can be described with the following algorithm:

- 1. Start with data on market shares of brands (both within their modules and within the wider group) in the factual (i.e. observed shares).
- Calculate the E matrix using equation (8), combined with equation (13) for Bertrand— (17) for Cournot—to obtain the initial L vector for Lerner indices.

- 3. Impose the new  $\Theta'$  ownership matrix implied by the merger.
- 4. Again using equation (8) followed by (13) for Bertrand or (17) for Cournot, calculate L', which enables us to compute the new markup ( $\mu = 1/(1 L)$ ):

$$\hat{\mu}_{bn} = \frac{1 - L_{bn}}{1 - L'_{bn}}.$$
(28)

5. Using equation (6) to compute the proportional change in brand b's share of group g expenditures in market n, we can then write

$$\hat{s}_{b|gn} = \hat{s}_{b|mn} \times \hat{s}_{m|gn} = \hat{s}_{b|mn} \left( \hat{P}_{mn}^M / \hat{P}_{gn}^G \right)^{1-\rho} = (\hat{\mu}_{bn} / \hat{\varphi}_{bn})^{1-\sigma} (\hat{P}_{mn}^M)^{\sigma-\rho} (\hat{P}_{gn}^G)^{\rho-1}, \quad (29)$$

where the last equality comes from substituting in the change in the market share of b within module m, that is given by

$$\hat{s}_{b|mn} = \left(\frac{\hat{\mu}_{bn}}{\hat{\varphi}_{bn}\hat{P}_{mn}^M}\right)^{1-\sigma}.$$
(30)

Since  $s_{bn} = s_{b|gn} \times s_{gn}$  and  $\hat{s}_{bn} = \hat{s}_{b|gn} (\hat{P}_{gn}^G / \hat{P}_n)^{1-\eta}$ , the change in the market share of b in all beverages is

$$\hat{s}_{bn} = (\hat{\mu}_{bn}/\hat{\varphi}_{bn})^{1-\sigma} (\hat{P}_{mn}^M)^{\sigma-\rho} (\hat{P}_{gn}^G)^{\rho-\eta} \hat{P}_n^{\eta-1}.$$
(31)

The change in the market share of brand *b* within group *g* is therefore a function of i) the change in markups from equation (28), ii) the observable change in brand *b*'s attributes  $\hat{\varphi}_{bn}$ , iii) observed initial market shares, and structural parameters. Indeed, with nested CES demand, the changes in the price indices have no other elements than the 3 above:

$$\hat{P}_{mn}^{M} = \left(\sum_{b} \mathbb{I}_{bn} s_{b|mn} (\hat{\mu}_{bn} / \hat{\varphi}_{bn})^{1-\sigma} \right)^{\frac{1}{1-\sigma}}.$$
(32)

$$\hat{P}_{gn}^{G} = \left(\sum_{m} s_{m|gn} (\hat{P}_{mn}^{M})^{1-\rho}\right)^{\frac{1}{1-\rho}},$$
(33)

$$\hat{P}_n = \left(s_{0n} + \sum_g s_{gn} (\hat{P}_{gn}^G)^{1-\eta}\right)^{\frac{1}{1-\eta}},\tag{34}$$

where  $s_{0n}$  is the share of the outside good in the economy, which is defined as the share of beverages other than beer and spirits in the entire beverage budget,  $X_n$ .

- 6. Update the Lerner index vector, L' by plugging the new market shares,  $s'_{bn}$ ,  $s'_{b|mn} = \hat{s}_{b|mn} s_{b|mn}$  and  $s'_{b|qn} = \hat{s}_{b|gn} s_{b|gn}$  into equation (8) to deliver a new E matrix.
- 7. Iterate steps 5 to 6 until convergence in the equilibrium market shares.

The outcomes of the counterfactual we examine are the changes in price indexes and in market concentration. The percentage change in the price index for each product category-market,  $\hat{P}_{gn} - 1$ , is described in equation (33). The counterfactual level of concentration is  $H'_{gn} = \sum_{f} (S'_{f'n})^2$ . A complete welfare calculation lies beyond the scope of this paper. This is because we do not know changes in fixed costs and, also, cannot map changes in profits to the nations of the ultimate claimants.⁴³

#### 5.2 Undoing forced divestitures: counterfactual results

Gutierrez and Philippon (2018) argue that the EU anti-trust authorities have been much more vigorous in preventing anti-competitive mergers than their US counterparts. In the beer industry, competition authorities on both sides of the Atlantic have forced divestitures to avoid concentration and even multi-market coordination effects.⁴⁴

AB InBev was compelled to divest large sets of brands in five separate cases. First, when InBev bought Anheuser Busch in 2008, it had to divest the US-market rights of Labatt brands (acquired in 1995) to a new company called North American Breweries (who later sold it to the Costa Rican firm FIFCO). Second, when it bought the Modelo Group, it had to divest the US-market rights of Corona and several other brands to Constellation Brands (a company mainly active in wine). The acquisition of SAB Miller in 2016 triggered forced divestitures in the US, EU, and China. Specifically, a package of popular EU brands was sold to Asahi, all the Miller brands were sold to MolsonCoors, and AB InBev's minority share of China Resources was sold to its Chinese partner.

Our model and data are well-suited to evaluate the efficacy of these divestitures by simulating a counterfactual in which the competition authorities permit AB InBev to retain all the brands it in fact had to divest. Specifically, we undo the divestitures described above and recompute the equilibrium in all markets. The results for the countries where the elimination of the divestiture is predicted to change the price index by more than

⁴³Multinational firms have complex capital structures and the rules of corporate taxation are equally difficult to apply on a global scale.

⁴⁴The ABI/Modelo decision by US DOJ and European Commission decision (Case M.7881: AB IN-BEV/SABMILLER) on the SABMiller acquisition points to both effects to justify divestitures.

one percent are displayed in Table 9. Sorted in descending order by the price change for Bertrand ( $\hat{\varphi}_{bn} = 1$ ), the table also includes prices changes for Cournot. The last two columns display the simulation results incorporating the adjustment to  $\varphi_{bn}$  predicted in our regression analysis for beer (the group fixed effects and HQ rows of the second column of coefficients in tables 7 and E.6).

Country	%Chg. <i>P_{gn}</i> Bertrand		%Chg. <i>P_{gn}</i> Bertrand	$\begin{array}{c} t \ (\hat{\varphi}_{bn} \neq 1) \\ \text{Cournot} \end{array}$
United States	3.60	4.25	3.25	3.85
Hungary	0.78	1.16	0.32	0.38
United Arab Emirates	0.67	0.95	0.72	0.96
Netherlands	0.64	1.23	0.21	0.69
Italy	0.62	1.00	0.33	0.57
Argentina	0.61	0.55	0.61	0.55
Australia	0.46	0.61	0.51	0.64
South Korea	0.44	0.53	0.49	0.56
Czechia	0.36	0.54	-0.59	-0.82
United Kingdom	0.35	0.57	0.31	0.50
Slovakia	0.11	0.18	-0.49	-0.81
Poland	0.00	0.00	-0.57	-0.90

Table 9: What if antitrust authorities had been more permissive?

Notes: The table reports the effect of undoing divestitures imposed by the US and the EU since 2007 on the percent change in the price index for beer in each country in 2018. To be included in this table, at least one absolute price change must exceed 1%.

The US consumer is by far the most important beneficiary of the forced divestitures. Had AB InBev been able to keep all the brands owned by the companies it acquired, the beer price index in the US would be about four percent higher. Cournot conduct increases the benefits of divestitures by two thirds of a percent. The third and fourth columns show that taking into account changes in  $\varphi_{bn}$  leads to a small reduction of the market power effects. The main reason is that AB InBev is considered to have dual headquarters in Belgium and New York. Hence, the non-divestiture to MolsonCoors (Miller) and Constellation Brands (Corona) does not change HQ frictions. Moreover, all the firms involved in the divestitures have the same or similar group fixed effects, except for FIFCO who obtained the relatively small Labatt.⁴⁵

The case of the United Arab Emirates (UAE) provides a clear example of the potential for positive spillovers in competition policy. The UAE did not force divestitures but

⁴⁵Non-divestiture to FIFCO helps (by very small amounts) in two ways: keeping Labatt with a better firm and keeping the headquarters in the US—rather than Costa Rica.

it benefited from the US and EU preventing AB InBev from keeping Miller and Peroni worldwide. It is a rare market where local stars are irrelevant; divestiture lowers the price index about a percent by promoting competition between global giants. The lead-ing brands are Heineken followed by four of AB InBev's global giants.

The EU commission's intervention protected consumers from increases in market power in Hungary, the Netherlands, and Italy that would have otherwise lead to a 0.5–2.0% increases in the price index. In Hungary, AB InBev keeps the Dreher Brewery local stars (accounting for 31% of the market) it had to divest to Asahi. This allows AB InBev to avoid competition for its global giants Stella Artois, Leffe, and Becks, which collectively held 7% of the Hungarian market. In Italy, AB InBev brands (led by Becks at 6%) accounted for 13% of the market in 2016, similar to Asahi's 14% (8% of which was Peroni). Cost increases (due to moving the HQ from Belgium to Japan) partially offset the market power effects.

The market situations in Slovakia and Poland exemplify the unintended consequences of divestitures to a remote owner. In these countries, the simulation predicts minimal (or zero in the case of Poland) price rises due to market power.⁴⁶ However, the move of HQ from Belgium to Japan increases frictions by enough to raise the price index of beer by 0.6 to 0.9%. The potential costs of distance between market and headquarters is an issue that can only be quantified by combining data from multiple markets.

In sum, the divestitures imposed by EU and US competition authorities reduced market power by enough to lower prices by half a percent to four percent in four countries relative to the permissive counterfactual. Unfortunately, in three countries, the replacement of a headquarters in nearby Belgium with one in Japan implies cost increases that more than offset the benefits. The mixed success of the actual remedies motivates the next set of policy counterfactuals, considering remedies that might have been applied.

### 5.3 Forcing counterfactual divestitures

Our second counterfactual examines whether competition agencies that were passive in response to AB InBev's acquisitions could have achieved net consumer savings by emulating the US/EU approach. The simulation reported in Table 10 reassigns the global rights for Labatt brands to FIFCO, the Modelo brands (including Corona) to Constellation, and all the local SABMiller brands to Asahi. Since FIFCO, Constellation, and Asahi had low or zero presence in the markets where these brands had high market shares, this

⁴⁶In Poland, AB InBev retained no other brands (above the GMID 0.1% threshold) after the divestiture. This implies no change in markups due to pure market power effects. The EU Commission justified the divestiture of the Polish brands due to concerns over multi-market contacts.

policy resembles placing the pricing decisions for these brands under independent control. The key difference is that the reallocation of ownership potentially changes headquarters frictions and firm effects.

Country			%Chg. P _{gn}	
	Bertrand	Cournot	Bertrand	Cournot
Colombia	-17.68	-15.62	-17.09	-14.52
Peru	-15.47	-11.01	-15.09	-10.19
Ecuador	-11.37	-9.68	-11.36	-9.04
Uruguay	-4.28	-4.28	-3.70	-3.56
Canada	-2.25	-3.62	-2.51	-4.03
Bolivia	-1.99	-3.17	-1.70	-2.66
Dominican Republic	-1.95	-1.26	-1.68	-1.03
Australia	-1.91	-3.39	-3.76	-4.70
Argentina	-1.56	-1.82	-1.52	-1.69
United Arab Emirates	-1.16	-1.76	-1.06	-1.42
Mexico	-1.13	-1.81	0.70	0.34
Chile	-0.93	-1.60	-0.52	-1.04
South Africa	-0.79	-1.08	-3.70	-3.45
India	-0.11	-0.30	-1.37	-1.46
Nigeria	-0.08	-0.15	-1.07	-0.88

Table 10: What if antitrust authorities had followed EU/US lead?

Notes: The table reports the effect of forcing divestitures on the percent change in the price index for beer in each country in 2018. To be included in this table, at least one absolute price change must exceed 1%.

The largest gains would accrue to consumers in three Andean countries where SAB-Miller had acquired the local star brands. Forcing divestitures would have reduced the beer price index by 9–18% depending on the country and assumptions. The Dominican Republic and Canada would also experience gains as large, or larger, than those generated by divestiture for the US.

Australia and Canada both issued no-action letters in 2016, commenting that they did not foresee adverse effects of the SABMiller acquisition on competition in their respective beer markets. Table 10 suggests that implementing the three divestitures (Labatt, Modelo, and SABMiller EU brands) would have saved Canadian consumers between 2.3% and 4%. Australian beer drinkers would gain 1.9% to 4.7%. Mexico could also have generated market power gains through compelling divestiture of the Modelo brands in the Mexican market. However, since our experiment divest towards US-based Constellation, the increased frictions overturn the overall result into negative territory.

The price reductions reported in Table 10 should be thought of as the cost-saving

for individual countries to deviate from their historical permissive behavior. Had every country insisted on divestiture, the acquisition itself would not make sense. To obtain consent for its purchase of SABMiller, AB InBev had to divest more than half of the 155 brands SABMiller offered in 2015. In 2019 they sold their Australian brand portfolio to Asahi. Taking into account all the subsequent brand divestitures, AB InBev paid a net price of \$83.4bn for the SABMiller brands it retained.⁴⁷ Our counterfactuals suggest the main benefit to AB InBev was near monopolization of several Latin American beer markets.

### 5.4 Restoring 2007 owners: counterfactual results

The final counterfactual can be framed as implementing a ban on all changes in brand ownership. The simulation calculates a new equilibrium using 2018 brand market shares as an input, but applying the 2007 mapping of brands to firms, that is o(b, 2007). The EHA procedure then calculates the counterfactual 2018 brand market shares.

Category	# of	Conduct	Chg. HHI		%Chg. P _{gnt}		
	Countries	assumed	Mean	Median	Mean	Median	
				with $\hat{\varphi}$	$\dot{\phi}_{bn} = 1$		
Beer	76	Bertrand	445	239	1.66	0.56	
Beer	76	Cournot	517	284	1.93	0.99	
Spirits	75	Bertrand	72	20	0.08	0.03	
Spirits	75	Cournot	74	26	0.18	0.07	
				with $\hat{\varphi}$	$\hat{\phi}_{bn} \neq 1$		
Beer	76	Bertrand	399	196	2.20	0.77	
Beer	76	Cournot	464	247	2.45	1.06	
Spirits	75	Bertrand	65	19	0.26	0.08	
Spirits	75	Cournot	66	21	0.37	0.13	

Table 11: Summary of outcomes of the counterfactual restoring 2007 brand owners

Notes: The table reports the mean and median change in the Herfindahl Index and in the percent change in the price index resulting from banning all ownership changes over the last 12 years (restoring 2007 owners). The bottom panel incorporates changes in brand type.

Table 11 summarizes counterfactuals run on 76 (beer) or 75 (spirits) markets. Ownership changes between 2007 and 2018 led to widespread increases in concentration. The US DOJ guidelines state that mergers in concentrated markets that raise the HHI by 200

⁴⁷The gross price paid in 2016 before any divestitures was \$122 billion. All values taken from *Financial Times*, "How deal for SABMiller left AB InBev with lasting hangover" (July 24, 2019).

points or more "will be presumed likely to enhance market power."⁴⁸ Table 11 points to mergers increasing market power by greater than the DOJ threshold in over half the beer markets. Compared to a counterfactual of no changes in ownership, the simulation points to price indexes that are 0.18–2.45% higher for the average country.⁴⁹ The biggest increase is for beer, assuming Cournot and including changes in brand type (i.e. the second row of the lower panel). The smallest changes are the pure market power effects of mergers in the spirits category (i.e. the third and fourth rows of the top panel).

Country	Chg.	HHI	%Chg. $P_{qnt}$		
-	Bertrand	Cournot	Bertrand	Cournot	
		Be	er		
Colombia	3513.13	4306.04	17.11	15.88	
Dominican Republic	2858.67	3115.12	25.59	22.21	
Ecuador	2666.51	3036.84	10.14	8.50	
Peru	1925.16	2567.80	14.18	9.64	
Nigeria	1352.17	1454.67	5.81	6.79	
		Spi	rits		
Tunisia	1116.40	1102.12	0.06	0.50	
Turkey	613.77	574.79	3.94	5.28	
Algeria	518.84	514.92	0.15	0.41	
Morocco	513.61	509.26	0.71	0.94	
Czechia	301.46	309.00	0.53	0.95	

Table 12: Changes in concentration in the most affected markets

Notes: The table reports the top 5 countries based on the increase of the HHI caused by mergers over the 2007-2018 period. Order starts by top increases in HHI for Bertrand conduct. All simulations incorporate changes in  $\hat{\varphi}_{bn}$ .

Table 12 reports the counterfactual concentration and price index changes for the most affected countries in our data set. The largest increases in concentration and biggest price index increases are for beer. The double-digit price rises for beer in South America are mainly attributable to the AB InBev acquisition of local stars that SABMiller had acquired previously. The Dominican Republic sees 22–25% price index increases coming from the local brand acquisition by AB InBev described in the introduction.

In the spirits industry, we see sizeable effects in just two markets: Turkey and Tunisia. In the former, Diageo's acquisition of the owner of Yeni Raki, the most popular spirit in

⁴⁸https://www.justice.gov/atr/horizontal-merger-guidelines-08192010

⁴⁹Most of the average price increases are smaller than 4% average that Kwoka (2014) obtained in his meta-analysis of retrospectives covering a variety of products.

the country, leads to a price rise of 4–5% (incorporating the higher costs from moving the HQ to London). The Tunisia case provides a rare example of market power rising entirely via the combination of global giant brands. Pernod-Ricard, whose Chivas and Ballantines brands had significant market shares (17% in 2018), bought the most popular spirit in Tunisia, Absolut (32% market share in 2018). This is a case where our nested structure has important consequences for the counterfactual. Without modules, the merger raises the Tunisian spirit price index by 3–4%. However, as Chivas and Ballantines are whiskys and Absolut is a Vodka, the three-layer structure dampens the effect of increased concentration in Tunisian spirits.

### 5.5 Rationalizing mergers and acquisitions

Our counterfactual method can predict the change in profits for each brand when they are combined under a new common owner. It is not enough to compute the change in profits for the owner since we must subtract the amount paid for acquired brands. To fix ideas, consider a simplified version of the purchase of Anheuser-Busch (A-B) by InBev, which led to the formation of AB InBev. The gain for InBev owners is the change in net profit for their brand, Stella, plus the net profit of Bud (A-B's brand) under InBev ownership *minus* the amount paid to A-B owners for Bud. The gain for A-B owners is the amount they receive for Bud minus the net profit of Bud under A-B ownership. For the merger to be mutually acceptable, both InBev and A-B owners have to benefit. Summing their gains, the acquisition price drops out, and we are left with the sum of the changes in profits net of fixed costs for Stella and Bud. Thus, acquisitions can only be rationalized for both parties if the sum of net profits increases.

Setting changes in fixed costs aside for the moment, the change in variable profits generated by the ownership changes is given by

$$\Delta \Pi_f = \sum_n \sum_{b \in \mathcal{F}_f} (\pi_{bn} - \pi_{bn}^0), \tag{35}$$

where  $\mathcal{F}_f$  is the set of brands that firm f owns after acquiring various companies with their brands (in the example  $\mathcal{F}_f$  is {Stella, Bud}). The variable profit of a brand under its true 2018 owner is denoted  $\pi_{bn}$ , whereas  $\pi_{bn}^0$  is the profit of the same brand if it were given back to its 2007 owner. In the absence of changes in appeal and costs, we would expect  $\Delta \Pi_f$  to be positive when brands are combined because the new owner could always retain the markups charged by the previous owners, but might be able to increase profits by raising markups on formerly competing products. There are two countervailing forces in our framework that could potentially lead to reductions in  $\varphi_{bn}$  and therefore variable profits: 1) frictions: appeal declines because a domestic brand becomes foreign headquartered (or headquartered further away), 2) costs rise because the new owner is less efficient (i.e. has a lower  $\varphi_f^F$ —as captured in the fixed effects—than the previous owner). Taking these two mechanisms into account we compute  $\Delta \Pi_f$  for each of the fthat acquired brands during the study period.

Turning to fixed costs, denote  $\Delta \Phi_f$  as the combined change in fixed costs for all the brands in  $\mathcal{F}_f$ . These combined fixed costs might fall due to economies of scope or better matches between brands and their owners. Alternatively, they might rise due to "span of control" effects (diseconomies of owning too many brands). To rationalize the merger, it must be that profits net of fixed costs do not fall, or  $\Delta \Phi_f \leq \Delta \Pi_f$ . Thus, if rising frictions reduce profits, then fixed costs must fall by enough to compensate. Thus,  $-\Delta \Pi_f$ , is the lower bound on the reduction in fixed costs for the major ownership changes observed in our study period.

The counterfactual we consider for this exercise is the same one we used in section 5.4. Thus, the merger gains we compute consider all the ownership changes over the 2007–2018 period simultaneously.⁵⁰

	$\hat{arphi}_{bn}$	= 1	$\hat{\varphi}_{bn} \neq 1$		
	Bertrand	Cournot	Bertrand	Cournot	
pct. firms with pos. gains	95.2	95.2	58.1	61.3	
pct. sales with pos. gains	99.9	99.8	61.5	71.2	

Table 13: Merger profitability summary

Notes: The table reports percentages for firms involved in at least one acquisition between 2007 and 2018. The first row gives the percent of firms with  $\Delta \Pi_f > 0$ , the second row gives their share of total sales.

Table 13 shows that the pure oligopoly effects of mergers are overwhelmingly positive: only one firm in 20 would be better off without the ownership changes that occurred between 2007 and 2018 and those losers account for negligible sales shares. The reason why a few firms lose in absence of any fix cost savings is that there were some important divestments in some markets. Aside from the ones covered in section 5, Diageo's purchase of the United Spirits brands in India resulted in the break-up of United Breweries, with the beer brands (notably Kingfisher) severed from the spirits brands (McDowells). In our nested structure this has the effect of reducing markups.

⁵⁰The alternative would have been to evaluate the mergers one by one. The method we chose builds in gains from complementarities between mergers.

Firm	$\Delta \Pi_f$ (with	$\mathbf{n}\hat{\varphi}_{bn}=1)$	$\Delta \Pi_f$ (with	$\hat{\varphi}_{bn} \neq 1$ )
	Bertrand	Cournot	Bertrand	Cournot
AB InBev	0.94	1.03	1.81	2.18
Heineken	0.60	0.86	-1.83	-2.09
Carlsberg	0.24	0.28	-0.03	0.05
MolsonCoors	0.70	1.01	2.27	2.90
Asahi	0.16	0.27	-0.71	-1.15
Kirin	0.62	1.03	-0.80	-0.93
Diageo	0.24	0.43	0.80	1.02
Pernod-Ricard	0.10	0.08	0.14	0.16
Suntory	0.11	0.20	-1.96	-2.33
Bacardi	0.17	0.31	-0.34	-0.35
Brown-Forman	0.17	0.33	0.68	0.88
Campari	0.20	0.33	0.52	0.67

Table 14: Merger profitability for the large players

Note: Counterfactual  $\Delta \Pi_f$  amounts reported as percentages of 2018 sales.

When we incorporate changes in cost-adjusted appeal, the situation is not nearly as rosy for mergers, but still the majority of firms benefit from the ownership changes. Digging deeper in Table 14, we see all the major acquiring firms benefited before inclusion of the  $\varphi_{bnt}$  reductions brought about by headquarter changes and differences in the owner effects. We see that AB InBev's gains amount to roughly two percent of 2018 sales. One reason they do so well is that their biggest purchases were SABMiller brands that were already foreign-headquartered. For Heineken, Carlsberg, Suntory, and Bacardi, the extra cost of foreign ownership largely wiped out or even turned negative their net gains in variable profits. Heineken and Suntory would need reductions in fixed costs equivalent to about two percent of sales to make their acquisitions pay off. This magnitude seems like an optimistic, but plausible, expectation of savings from closing duplicative breweries or centralizing marketing and distribution.

## 6 Conclusion

In the beer and spirits industries, a small group of large firms, headquartered in a handful of countries, has expanded primarily via cross-border acquisitions. This process of multinational brand amalgamation has the potential to impact competition in a number of different ways. On the efficiency side, merging firms have long justified horizontal combinations on the basis of synergies. Competition authorities, on the other hand, have at times rejected mergers that were predicted to harm consumers. This paper obtains several new findings related to this debate. First, we find that brand type—extracted from data on market shares—is, for the most part, invariant to the identity of the owner. That is, after mitigating limited mobility bias, firm fixed effects explain 4% or less of the variation in a brand's cost-adjusted appeal.

There is one way that ownership *does* affect cost-adjusted appeal, however. In the spirits industry, and to a lesser extent, in the beer industry, we estimate that brand type is higher in the countries where their owners are headquartered. Our results imply an 11–12% penalty on cost-adjusted appeal from foreign acquisitions with little in the way of predictable efficiencies. Nevertheless, our counterfactuals imply rises in market power led to increased profits for the majority of brand-acquiring firms. The changes in market power varied greatly across markets. There is a simple heuristic for identifying cases where M&A is harmful: Consumer surplus falls the most when foreign firms owning global giant brands acquire the domestic owners of local star brands in the same market segment.

Cross-country comparisons in our counterfactuals quantify the beneficial role of competition policy towards mergers. Divestitures forced by the US and EU led to significant consumer savings. Canada and Australia could have achieved similar savings by imposing divestitures along the same lines. The greatest potential for the use of these structural remedies would be in Colombia, Ecuador and Peru, where counterfactuals reveal that beer price increases of 9 to 18% could have been avoided. The more permissive competition authorities allowed Europe-based multinationals to extract a substantial part of potential consumer surplus from lower income countries.

We conclude with a caution against the indiscriminate application of lessons drawn from the analysis of beer and spirits mergers to other sectors. Obviously, research and development is much more important in electronics, software, and pharma industries. Nothing in this paper can indicate how cross-border acquisitions affect innovation. Nevertheless, in sectors as diverse as dog food, eyeglasses, and chocolate bars, the GMID data exhibit similar patterns of multinational brand amalgamation. Hence, we believe the issues we raise here—and the methods we have employed—have potentially broad applications.

## References

Abowd, J. M., F. Kramarz, and D. N. Margolis (1999). High wage workers and high wage firms. *Econometrica* 67(2), 251–333.

- Andrews, M. J., L. Gill, T. Schank, and R. Upward (2008). High wage workers and low wage firms: negative assortative matching or limited mobility bias? *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 171(3), 673–697.
- Ashenfelter, O. and D. Hosken (2010). The effect of mergers on consumer prices: Evidence from five mergers on the enforcement margin. *The Journal of Law and Economics* 53(3), 417–466.
- Ashenfelter, O., D. Hosken, and M. Weinberg (2014). Did robert bork understate the competitive impact of mergers? evidence from consummated mergers. *The Journal of Law and Economics* 57(S3), S67–S100.
- 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.
- Atkeson, A. and A. Burstein (2008). Pricing-to-market, trade costs, and international relative prices. *The American Economic Review* 98(5), 1998–2031.
- Autor, D., D. Dorn, L. F. Katz, C. Patterson, and J. Van Reenen (2020). The fall of the labor share and the rise of superstar firms. *The Quarterly Journal of Economics*, ??–??
- Bernard, A. B., J. B. Jensen, S. J. Redding, and P. K. Schott (2007). Firms in international trade. *Journal of Economic Perspectives* 21(3), 105–130.
- 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. T. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, 242–262.
- Björnerstedt, J. and F. Verboven (2016). Does merger simulation work? Evidence from the Swedish analgesics market. *American Economic Journal: Applied Economics* 8(3), 125–64.
- Blonigen, B. A. and J. R. Pierce (2016). Evidence for the effects of mergers on market power and efficiency. Working Paper 22750, National Bureau of Economic Research.
- Bonhomme, S., K. Holzheu, T. Lamadon, E. Manresa, M. Mogstad, and B. Setzler (2023). How much should we trust estimates of firm effects and worker sorting? *Journal of Labor Economics* 41(2), 291–322.

- Bonhomme, S., T. Lamadon, and E. Manresa (2019). A distributional framework for matched employer employee data. *Econometrica* 87(3), 699–739.
- Bonhomme, S. and E. Manresa (2015). Grouped patterns of heterogeneity in panel data. *Econometrica* 83(3), 1147–1184.
- Brander, J. A. and P. R. Krugman (1983). A 'reciprocal dumping' model of international trade. *Journal of International Economics* 15(3), 313–321.
- Burstein, A., V. Carvalho, and B. Grassi (2019). Bottom-up markup fluctuations. mimeo.
- Card, D., J. Heining, and P. Kline (2013). Workplace Heterogeneity and the Rise of West German Wage Inequality. *The Quarterly Journal of Economics* 128(3), 967–1015.
- Chung, F. R. (1997). Spectral graph theory. Number 92. American Mathematical Society.
- Coşar, K., P. Grieco, S. Li, and F. Tintelnot (2018). What drives home market advantage? *Journal of International Economics* 110, 135–150.
- Conlon, C. T. and N. Rao (2015). The price of liquor is too damn high: Alcohol taxation and market structure. *NYU Wagner Research Paper* (2610118).
- Costinot, A. and A. Rodriguez-Clare (2014). Trade theory with numbers: Quantifying the consequences of globalization. In E. Helpman (Ed.), *Handbook of International Economics*, Volume 4. Elsevier.
- Cunningham, C., F. Ederer, and S. Ma (2019). Killer acquisitions. *Mimeo*.
- Dafny, L., M. Duggan, and S. Ramanarayanan (2012). Paying a premium on your premium? Consolidation in the US health insurance industry. *American Economic Review* 102(2), 1161–85.
- De Loecker, J. and J. Eeckhout (2018). Global market power. Working Paper 24768, National Bureau of Economic Research.
- De Loecker, J., J. Eeckhout, and G. Unger (2020). The Rise of Market Power and the Macroeconomic Implications. *The Quarterly Journal of Economics*.
- 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 (2016). Estimating market power: Evidence from the US brewing industry. Working Paper 22957, National Bureau of Economic Research.

- De Loecker, J. and F. Warzynski (2012). Markups and firm-level export status. *American Economic Review* 102(6), 2437–2471.
- Dekle, R., J. Eaton, and S. Kortum (2008). Global rebalancing with gravity: Measuring the burden of adjustment. *IMF Staff Papers* 55(3), 511–540.
- Dinlersoz, E. M. and M. Yorukoglu (2012). Information and industry dynamics. *The American Economic Review* 102(2), 884–913.
- Edmond, C., V. Midrigan, and D. Y. Xu (2015). Competition, markups, and the gains from international trade. *The American Economic Review* 105(10), 3183–3221.
- European Commission (2008). Pernod Ricard/ V&S. Case No COMP/M.5114.
- Gaubert, C. and O. Itskhoki (2018). Granular comparative advantage. Working Paper 24807, National Bureau of Economic Research.
- Goldberg, P. K. and F. Verboven (2001). The evolution of price dispersion in the European car market. *The Review of Economic Studies* 68(4), 811–848.
- Gourio, F. and L. Rudanko (2014). Customer capital. *Review of Economic Studies 81*(3), 1102–1136.
- Grullon, G., Y. Larkin, and R. Michaely (2019). Are US industries becoming more concentrated? *Review of Finance* 23(4), 697–743.
- Gutierrez, G. and T. Philippon (2018). How EU markets became more competitive than US markets: A study of institutional drift. Working Paper 24700, National Bureau of Economic Research.
- Hausman, J., G. Leonard, and J. D. Zona (1994). Competitive analysis with differenciated products. *Annales d'Economie et de Statistique*, 159–180.
- Head, K. and T. Mayer (2019). Brands in motion: How frictions shape multinational production. *American Economic Review* 109(9), 3073–3124.
- Head, K. and T. Mayer (2021). Poor substitutes? Counterfactual methods in IO and trade compared. Discussion Paper 16762, Centre for Economic Policy Research. mimeo.
- Hottman, C. J., S. J. Redding, and D. E. Weinstein (2016). Quantifying the sources of firm heterogeneity. *The Quarterly Journal of Economics* 131(3), 1291–1364.

- Jochmans, K. and M. Weidner (2019). Fixed-effect regressions on network data. *Econometrica* 87(5), 1543–1560.
- Khandelwal, A. K., P. K. Schott, and S.-J. Wei (2013). Trade liberalization and embedded institutional reform: Evidence from Chinese exporters. *American Economic Review* 103(6), 2169–95.
- Kline, P., R. Saggio, and M. Sølvsten (2020). Leave-out estimation of variance components. *Econometrica*.
- Kwoka, J. (2014). *Mergers, merger control, and remedies: A retrospective analysis of US Policy*. MIT Press.
- Mayer, T. and G. Ottaviano (2007). *The Happy Few: The Internationalisation of European firms*. Bruegel Blueprint Series.
- Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica* 71(6), 1695–1725.
- Miller, N. H., G. Sheu, and M. C. Weinberg (2021, October). Oligopolistic price leadership and mergers: The United States beer industry. *American Economic Review* 111(10), 3123– 59.
- Miller, N. H. and C. M. Weinberg (2017). Understanding the price effects of the Miller-Coors joint venture. *Econometrica* 85(6), 1763–1791.
- Miravete, E. J., K. Seim, and J. Thurk (2018). Market power and the Laffer curve? mimeo.
- Nocke, V. and N. Schutz (2018). Multiproduct-firm oligopoly: An aggregative games approach. *Econometrica* 86(2), 523–557.
- Pinkse, J. and M. E. Slade (2004). Mergers, brand competition, and the price of a pint. *European Economic Review* 48(3), 617–643.
- Redding, S. J. and D. E. Weinstein (2018). Accounting for trade patterns. Discussion Paper 12446, Center for Economic Policy Research.
- Strang, G. (2005). *Linear algebra and its applications, 4th ed.* Brooks/Cole.
- 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.
- Verboven, F. (1996). International price discrimination in the european car market. *The RAND Journal of Economics* 27(2), 240–268.
- World Bank, T. (2020). World Development Report 2020: Trading for development in the age of global value chains. World Bank Publications.

# Appendix

# A Modules

Since our original data from GMID does not classify brands into greater detail than beer and spirits, we have enlisted several sources of this information. First, some brands (e.g. Seagram's Gin and Gin Lubuski) have their type revealed as part of the brand name. This also helps us identify low-alcohol and low-calorie beers. Second, we used a definition of modules similar to that employed by Nielsen's Homescan and the Iowa Liquor Control Board. Third, we aggregated detailed beer "styles" provided by the online rating site ratebeer.com into Nielsen-level modules. So doing, we have classified 4908 brands into modules, about 85% of the beer and spirits brands in GMID, and 97% of the sales for classifiable brands.⁵¹

## **B** Extensive margins for brands and markets

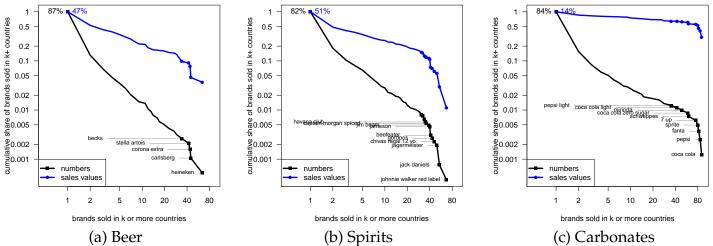
In this section, we document the very important cross-sectional extensive margin of market entry as well as the relatively small entry rates over time for beer and spirits.

Figure B.1 illustrates a few features of the distribution of brands across markets that play important roles in determining the outcomes of brand ownership changes in the beer and spirits industries. First, echoing a result shown repeatedly for exporters, a "happy few" brands are offered in many destinations and account for a disproportionate share of sales.⁵²

⁵¹The brands GMID lumps together as "Others" could only be assigned to modules by assumption.

⁵²Bernard et al. (2007) show these patterns in US data, Mayer and Ottaviano (2007) coin the term and show that the pattern holds for many countries.

Figure B.1: Global giants are rare



Notes: Symbols mark brands sold in > 30 countries. Log scales on both axes. Calculations exclude fringe brands since their counts and destinations are not known.

Sample	Add rate	Drop rate									
frame	(in percent)	(in percent)									
	Beer										
(a) Brand-level births and deaths:											
All brand/years	3.50	2.54									
Brands changing owners: before	NA	2.44									
Brands changing owners: after	NA	2.95									
(b) Brands added/dropped in a n	narket:										
All brand/market/years	0.06	2.63									
Continuing brands	0.03	0.76									
Brands changing owners: before	0.03	0.60									
Brands changing owners: after	0.03	1.34									
Spirits											
(a) Brand-level births and deaths	5:										
All brand/years	2.33	1.98									
Brands changing owners: before	NA	2.09									
Brands changing owners: after	NA	1.62									
(b) Brands added/dropped in a n	narket:										
All brand/market/years	0.06	1.85									
Continuing brands	0.03	0.72									
Brands changing owners: before	0.04	0.89									
Brands changing owners: after	0.04	1.50									

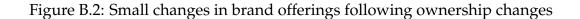
Table B.1: Adding and dropping brands in markets and overall: Beer and Spirits

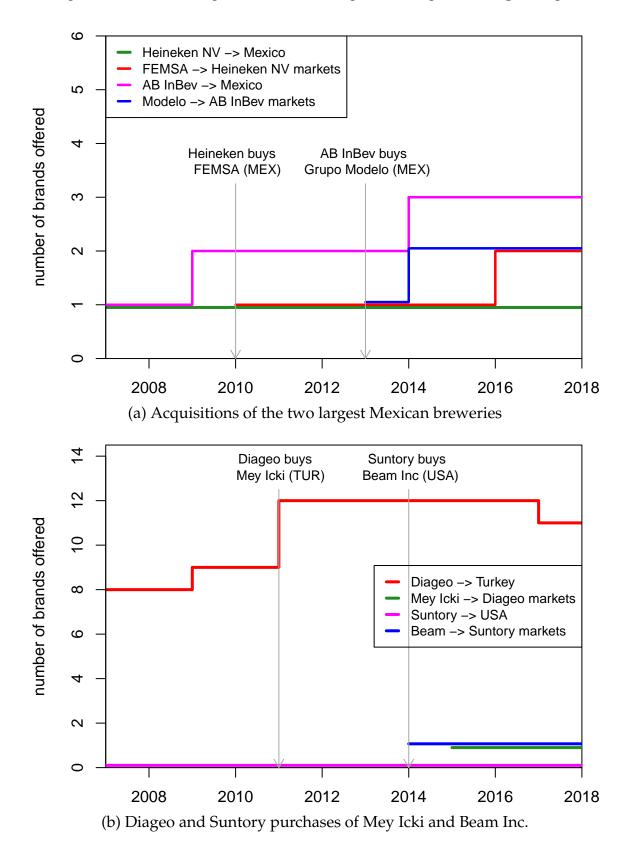
Table B.1 investigates the entry margin, through which firms add or drop brands in selected markets or altogether. The first panel considers the fraction of brands that are new each year (the add rate) whereas the second column is the fraction of brands that existed in the previous year but not the current year. Add rates are slightly higher (2.2 and 3.5%) than drop rates (1.6–3%). The drop rate does not fall in beer after acquisition and it does not fall much for spirits. Rather than the "buy to kill" pattern observed by Cunningham et al. (2019) in the pharmaceutical industry, firms in the beer and spirits industries "buy to keep." This difference is just what industrial organization would predict. While it can make sense to drop products in their early stages to save on development costs, most beer and spirits brands are already established in their markets. Therefore it makes more sense to simply raise their prices than to drop them. Note that add rates are not formulated in a way that would allow us to compare them before and and after acquisitions.

Panel (b) of table B.1 calculates add rates as a fraction of the number of potential market-years where the brand is absent in the previous period. The add rates are so small because there are 78 countries where brands might be offered but the vast majority are sold at home only. The second column shows the rate at which brands exit markets. Here the denominator is much smaller. Nevertheless, only two to three percent of brands are dropped from a market each year. Most of those exiting brands disappear because the brand itself was dropped. Among continuing brands, the exit rate is less than one percent. There is a slight uptick after acquisitions but over 98% of brand-market combinations are retained on a year-by-year basis.

Overall, we see high stability over time in which brands are offered and where they exceed the 0.1% market share threshold. Furthermore, changes in ownership do not seem to spur significant elimination of brands. Nor do they spur increased distribution across markets. This last result might seem surprising given the importance of global giants. It is based on the whole sample and might hide interesting dynamics for the big players. We therefore consider two case studies that demonstrate the limited extensive margin exhibited even by major acquisitions carried out by the largest firms in each industry.

Figure B.2(a) displays the temporal relationship between brand offerings in the buyer and target markets before and after two acquisitions of large Mexican beer makers. Before Heineken purchased FEMSA, it already sold Heineken in Mexico. Similarly AB InBev already offered Budweiser and Bud Light. After the 2010 and 2013 takeovers, Heineken did not bring any of its 302 brands to Mexico and AB InBev brought only its Belgian flagship brand, Stella Artois. In the reverse direction, Heineken ultimately put two of FEMSA's 14 brands in markets FEMSA did not previously serve. AB InBev put two of Grupo Modelo's 13 brands in a total of four new markets by 2018 (out of a possible 73





markets).

Figure B.2(b) examines two similar cases from the spirits category. Again we see very little in the way of expansion along the extensive margin following the acquisition of the Turkish Mey Icki, by Diageo, and of the acquisition of the American company Beam Inc. by Suntory. Diageo, owner of 204 brands, added just three new brands in Turkey (though it later dropped one) and took Mey Icki's top brand, Yeni Raki, to Bulgaria only (though it could potentially have offered it in 73 countries). None of Suntory's 63 brands had sales in the US that are large enough to cross the 0.1% GMID threshold—before or after the purchase of Beam.

These case studies focus on acquisitions which took place sufficiently long ago to observe their consequences. They show very small changes in brand offerings relative to the sizes of the firms involved. The case study results are consistent with the absence of noticeable changes in the rate of adding brands to markets, seen in table B.1.

## C Endogenous mobility bias: a quantification

Here we investigate the direction and size of bias from assignment processes that depart from equation (27). In particular, we specify an assignment process we call *idiosyncratic sorting* in which brand *b* is more likely to be assigned to firm *f* with which they have strong bilateral affinity, denoted  $\xi_{bf}$ . In the proposed data generating process,  $\xi_{bf}$  enters the determination of  $\varphi_{bnt}$  and also influences the brand acquisition decision.

The actual assignment process observed in the beer and spirits industries features multi-brand firms acquiring and absorbing other multi-brand firms. We cannot do justice to the complexities of this process here, which we view as the subject for a separate paper. Instead we model a stylized assignment process that captures the key economic principles and their econometric implications. In our DGP, N brands are assigned to N firms in year t based on the value generated by each brand-firm combination:

$$v_{bft} = \varphi_b^B \varphi_f^F \exp(\xi_{bf}) - \Phi_{bft}, \tag{36}$$

where the first term models variable profits as being multiplicative in the brand and firm level determinants of cost-adjusted quality ( $\varphi_b^B$  and  $\varphi_f^F$ ) and the idiosyncratic quality of the match ( $\xi_{bf}$ ). The last term,  $\Phi_{bft}$  represents the fixed costs incurred by firm f when it produces and sells the products of brand b. This term is important for two reasons. First, it is needed to generate mobility of brands across firms over time. Second, it introduces a random component to assignment that has no effect on the observed cost-adjusted appeal. Replacing  $\exp(\xi_{bf})$  with its expectation in equation (36) leads to an assignment process that satisfies equation (27). We will refer to this case as *hierarchical sorting* since assignment depends only on the ordering of  $\varphi_b^B$  and  $\varphi_f^F$  (and chance via the  $\Phi_{bft}$  shocks).

Instead of modeling the process of buying and selling brands, we assume that a candidate assignment matrix  $\Omega$  should have the feature that there are no mutually profitable reassignments. We do this by selecting the  $\Omega$  that maximizes industry profits. The equilibrium assignment in each period is the  $\Omega_{bft}$  that solves the linear program

Maximize 
$$\sum_{f=1}^{N} \sum_{b=1}^{n} v_{bft} \Omega_{bft}$$
, subject to  $\sum_{b=1}^{N} \Omega_{bft} = 1$ ,  $\sum_{f=1}^{N} \Omega_{bft} = 1$ ,  $0 < \Omega_{bft} < 1$ 

The first constraint ensures that each brand is assigned to a firm and the second constraint implies that all firms have a brand. The solution to this problem always respects  $\Omega_{bft} \in \{0,1\}$ . We solve for a new  $\Omega$  matrix in each period t, with brands potentially changing owners based on realizations of  $\Phi_{bft}$ .

To implement this DGP, we set  $\varphi_b^B = b$  for  $b = 1 \cdots N$ ,  $\varphi_f^F = f$  for  $f = 1 \cdots N$ , with N = 100. The idiosyncratic matching term,  $\xi_{bf}$ , is distributed Normal(0,1). On average, brands move to firms with whom they have good fit, which implies that  $\xi_{bf}$  in the selected sample has an expectation greater than zero. Fixed costs are  $\Phi_{bft} \sim \text{LogNormal}(8,1)$ . The solution of the model repeats for T periods. As T increases, firms connect to each other via brands that have been held in common. Furthermore, within the largest connected set, the connectivity index  $\lambda_2$  rises.

Figure C.1 displays the contribution of the firm (owner) fixed effect ( $\varphi^F$ ) to explaining the variance in  $\varphi_{bnt}$  in equation 25. The blue lines correspond to a DGP that allows for assortative matching but rules out matching based on *bf* affinity. We see the familiar limited mobility bias (LMB) result of Andrews et al. (2008) and Bonhomme et al. (2019) that firm shares are overestimated. As mobility increases, the estimated share converges to the true share (which is almost flat because the numerator is not random and the denominator is stable because of the law of large numbers). In contrast, the red lines illustrate endogenous mobility bias (EMB) coupled with LMB. Both biases are upward but only the LMB disappears through increases in the number of periods. One can decompose the total bias into the LMB component—the gap between the last red circle at T = 1000 and the horizontal dashed line—and the EMB—the gap between the red circle and the nearly horizontal solid red line.

Figure C.2 shows the results from Andrews et al. (2008) and Bonhomme et al. (2023) that correlations between firm and brand fixed effects are biased downwards by limited

Figure C.1: Firm (owner) shares in the variance of brand performance are biased upwards by both limited and endogenous mobility

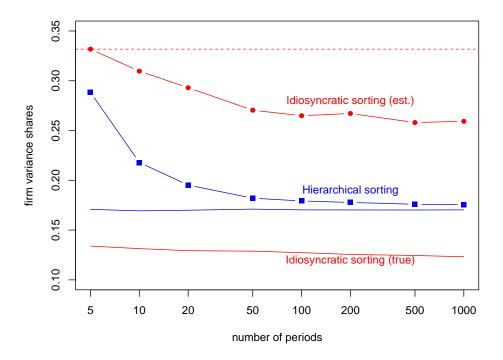
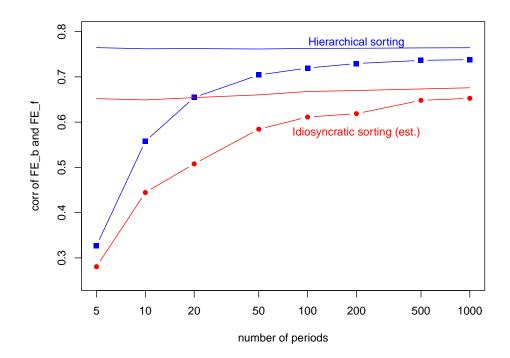
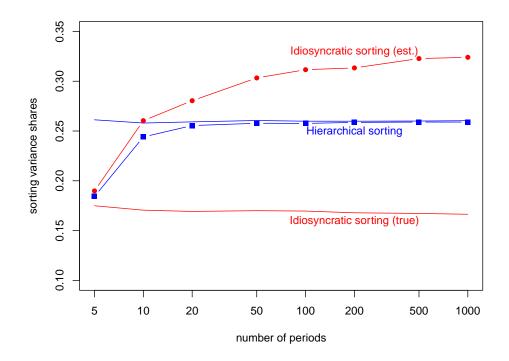


Figure C.2: The correlation between brand and firm fixed effects is biased downwards, but only due to limited mobility



mobility bias, which disappears for Hierarchical sorting (the AKM assumptions) as the number of periods increase. This is seen in the blue square line converging to the flat solid blue line. Somewhat surprisingly the bias in the correlation also gradually disappears as the number of periods grows for Idiosyncratic sorting that violates the AKM orthogonality condition. However, this requires unrealistically large numbers of periods to eliminate the limited mobility bias. Why is there little or no bias coming purely from endogenous mobility? In figure C.3 we see in the red lines that Idiosyncratic sorting does indeed bias the covariance upwards. So it appears that the lack of bias in the correlations comes from countervailing effects in the numerator and denominator of the correlation formula.

Figure C.3: The sorting shares  $(2 \times \text{ cov})$  in the variance of brand performance are biased upwards by endogenous mobility and downwards by limited mobility



The bottom line we draw from this investigation is that the role of firms in explaining variance in  $\varphi_{bn}$  is biased upwards by both limited mobility and endogenous mobility. However, once we take steps in our econometrics of the main text to mitigate limited mobility bias, the estimated firm shares are very small. Hence, bias coming from endogenous mobility should be very small as well. This is corroborated by the event study evidence in figure 5 and the low explanatory power of the brand-firm interactions.

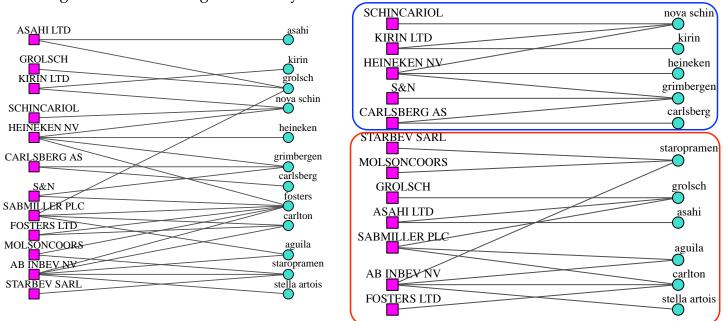
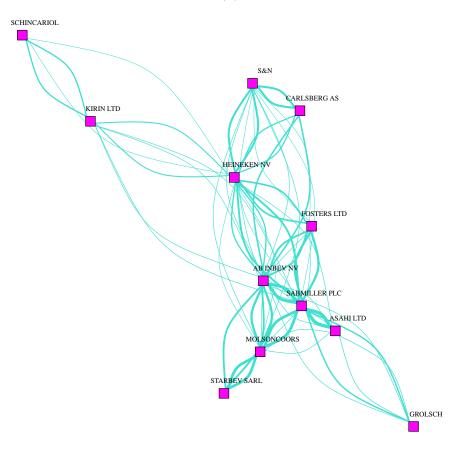


Figure D.1: Visualizing connectivity via an illustrative subset of brands and firms

(a) A connected set of firms and brands

(b) Without the Fosters brand, the sets disconnect



(c) The *induced* firm-to-firm network from panel (a)

Product group	# Firms		Mobility		Sales share	
Beer	90	21	13.6	50.6	80.0	70.8
Spirits	93	18	8.3	32.6	57.5	42.0
$\geq 10$ movers		$\checkmark$		$\checkmark$		$\checkmark$

Table D.1: Brand mobility in the largest connected set

Notes: # Firms is the count of firms in the largest connected set with and without the restriction of 10 or more moving brands per firm. Mobility is the average number of ownership changes per firm in the specified set. Sales share is the set's percentage of world sales.

## **D** Connectivity of the brand-firm network

In the third and fourth columns of Table D.1, we report the mobility ratios for all beverages, showing it for the largest connected set, and within that group, for the firms that experience more than ten movements (the large mobility group). Beer, and to a slightly lesser extent spirits, are characterized by two desirable features in this type of regressions: a high number of ownership changes, combined with a large share of world sales accounted for by firms in the connected set (shown in columns 5 and 6).

Figure D.1 illustrates the near-disconnectedness problem with an illustrative subset of firms and brands. Without the Fosters brand, the upper section of this graph (Schincariol, Kirin, Scottish & Newcastle, Carlsberg, and Heineken) would detach itself from the rest. While in this example Fosters is a "bottleneck" brand in the terminology of Kline et al. (2020), in the full dataset it can be removed without disconnecting Carlsberg, Heineken, and Kirin from AB InBev. The KSS leave-one-out set of firms comprises all the major beer makers.

Chung (1997) showed how the eigenvectors of the graph capture whether network is just connected or thickly connected. Jochmans and Weidner (2019) Theorem 2 shows that higher connectivity of the network, measured by  $\lambda_2$ , shrinks the upper bound for the variance of the estimates of the fixed effects (of firms). In a bipartite network, edges connect two sets of nodes where the only connections are between nodes from different sets. There is an *induced firm-to-firm network* with weighted edges between firms. The edge weight w(u, v) is an increasing function of in-common brand-market-years, with zero weight of a node to itself (w(u, u) = 0). The *Laplacian* of the weighted firm-to-firm network is a matrix with L(u, v) = -w(u, v) and  $L(u, u) = d_u$ , where  $d_v = \sum_u w(u, v)$ . In the case where w = 1,  $d_v$  is the degree, that is the number of edges connecting to vertex v. The elements of the *normalized* Laplacian are given by  $\mathcal{L}(u, v) = -w(u, v)/\sqrt{d_u d_v}$ and  $\mathcal{L}(u, u) = 1$ . As the smallest eigenvalue of each connected network is always zero, we refer to the smallest *positive* eigenvalue of  $\mathcal{L}$  as  $\lambda_2$ . Chung (1997) shows that the maximum  $\lambda_2$  in an unweighted network is n/(n-1), which occurs when each node has an edge to every other node. As the number of nodes grows large,  $\lambda_2 \rightarrow 1$ .

For all  $u \neq v$ , Jochmans and Weidner (2019) specify the weights as

$$w(u,v) = \sum_{b} \frac{n_{ub} n_{vb}}{N_b},$$

where  $n_{ub}$  is the count of market-years where brand *b* belongs to firm *u* and

$$n_{ub} = \sum_{nt} \mathbf{1}_{b \in \mathcal{F}_u} \times \mathbf{1}_{s_{bnt} > 0},$$

and  $N_b$  is the brand's total market-years under all owners:

$$N_b = \sum_f n_{fb}.$$

Figure D.1(c) shows the induced network of firm-to-firm links where the turquoise edges are based on brand-market-years. The thickness of these lines is proportional to the log of the Jochmans and Weidner (2019) weights described above. In this panel, *all* the brands are used in the weight calculation, not just the 12 illustrative brands in panel (a).

# **E** Additional regression results

		-		
			Bertrand	Cournot
	$\ln s_{bn}$	$\ln A_{bn}$	$\ln \varphi_{bn}$	$\ln \varphi_{bn}$
home	$1.155^{a}$	$0.171^{a}$	$0.314^{a}$	$0.328^{a}$
	(0.124)	(0.051)	(0.033)	(0.034)
distance	$-0.245^{a}$	$-0.036^{c}$	$-0.063^{a}$	$-0.064^{a}$
	(0.042)	(0.019)	(0.011)	(0.011)
common language	$0.216^{b}$	0.029	$0.057^{a}$	$0.059^{a}$
0 0	(0.085)	(0.038)	(0.022)	(0.023)
home (HQ)	$0.287^{a}$	$0.077^{b}$	$0.087^{a}$	$0.095^{a}$
	(0.100)	(0.038)	(0.027)	(0.027)
distance (HQ)	0.034	$0.022^{c}$	0.008	0.007
	(0.029)	(0.011)	(0.008)	(0.008)
com. lang. (HQ)	0.091	0.002	0.023	0.024
0	(0.060)	(0.025)	(0.017)	(0.017)
Observations	95,299	95,299	95,299	95,299
R ²	0.813	0.719	0.649	0.653

Table E.1: Pooled beer + spirits regressions, without firm fixed effects

Notes: Standard errors in (), clustered by origin-market dyads. Fixed effects at the brand-product and market-year-product dimensions included in each specification. HQ variables defined with respect to brand owner's headquarters country. Significance levels: 1% (*a*), 5% (*b*), and 10% (*c*).

	1 0			0
			Bertrand	Cournot
	$\ln s_{bn}$	$\ln A_{bn}$	$\ln \varphi_{bn}$	$\ln \varphi_{bn}$
home	$1.230^{a}$	$0.208^{a}$	$0.331^{a}$	$0.344^{a}$
	(0.150)	(0.058)	(0.040)	(0.041)
distance	$-0.225^{a}$	-0.021	$-0.058^{a}$	$-0.059^{a}$
	(0.047)	(0.021)	(0.012)	(0.013)
common language	$0.217^{b}$	0.032	$0.058^{b}$	$0.059^{b}$
	(0.095)	(0.043)	(0.025)	(0.025)
home (HQ)	$0.261^{c}$	0.047	$0.094^{a}$	$0.110^{a}$
	(0.133)	(0.048)	(0.035)	(0.036)
distance (HQ)	0.056	$0.024^{c}$	0.015	0.014
	(0.039)	(0.015)	(0.010)	(0.010)
com. lang. (HQ)	0.113	0.005	0.032	$0.034^c$
	(0.072)	(0.029)	(0.020)	(0.020)
Observations	65,061	65,061	65,061	65,061
$\frac{R^2}{N}$	0.790	0.675	0.611	0.615

Table E.2: Pooled beer + spirits regressions within the largest connected set

Notes: The sample is restricted to the largest connected set, within a product category. Standard errors in (), clustered by origin-market dyads. Fixed effects at the brand-product, firm, and market-year-product dimensions included in each specification. HQ variables defined with respect to brand owner's headquarters country. Significance levels: 1% (*a*), 5% (*b*), and 10% (*c*).

### E.1 Correlations of brand and firm fixed effects, with low mobility bias

Here we show the full set of correlation and variance shares for the fixed effects obtained in four different regressions using market shares, appeal, and cost-adjusted appeal (calculated under both conduct assumptions) as the dependent variables.

Table E.3 shows fixed effect correlations for regressions on all firms in the largest connected set. The underlying regressions in table E.4 apply the AGSU restrictions (keeping only moving brands and high mobility firms) to the estimating sample. In each table, the diagonal shows the ratio of the variance of the relevant fixed effect to the variance of the dependent variable. The off-diagonal elements of Table E.4 show the sign and magnitude of assortative matching.

	Brand					Fi	rm	
Dep. var.:	share	appeal	type B	type C	share	appeal	type B	type C
1	$(s_{bn})$	$(A_{bn})$	$(\varphi_{bn})$	$(\varphi_{bn})$	$(s_{bn})$	$(A_{bn})$	$(\varphi_{bn})$	$(\varphi_{bn})$
Beer	(,	(,	(,,	(,,	(,	(,	(,,	(1 )
brand market share	0.898							
brand appeal	0.761	0.912						
brand type B	0.993	0.746	0.889					
brand type C	0.988	0.742	0.999	0.876				
firm market share	-0.320	-0.210	-0.319	-0.314	0.118			
firm appeal	-0.253	-0.306	-0.243	-0.238	0.787	0.171		
firm type B	-0.304	-0.195	-0.308	-0.304	0.986	0.762	0.118	
firm type C	-0.291	-0.185	-0.295	-0.292	0.974	0.750	0.996	0.115
Spirits								
brand market share	0.891							
brand appeal	0.689	0.867						
brand type B	0.998	0.689	0.889					
brand type C	0.997	0.687	1.000	0.886				
firm market share	-0.456	-0.239	-0.452	-0.452	0.357			
firm appeal	-0.377	-0.372	-0.377	-0.378	0.764	0.229		
firm type B	-0.453	-0.240	-0.450	-0.451	0.996	0.766	0.365	
firm type C	-0.450	-0.239	-0.448	-0.449	0.993	0.767	0.999	0.373

Table E.3: Correlations between fixed effects in the largest connected set

Notes: For brand and firm type, we use B and C to denote Bertrand and Cournot conduct, respectively. **Diagonal:** ratio of FE variances to variance of the dependent variable. **Off-diagonal:** correlation. Underlying regressions keep the largest connected set.

As found in AGSU, the patterns of correlation in the largest connected set exhibit *negative* assortative matching: all correlations between brands and firm fixed effects are negative and large in absolute value, for both beer and spirits. After imposing the AGSU restrictions in Table E.4, the correlations become much smaller, and not even systemat-

	Brand					Fi	rm	
Dep. var.:	share	appeal	type B	type C	share	appeal	type B	type C
	$(s_{bn})$	$(A_{bn})$	$(\varphi_{bn})$	$(\varphi_{bn})$	$(s_{bn})$	$(A_{bn})$	$(\varphi_{bn})$	$(\varphi_{bn})$
Beer								
brand market share	0.831							
brand appeal	0.798	0.872						
brand type B	0.994	0.790	0.824					
brand type C	0.990	0.785	0.999	0.819				
firm market share	-0.126	-0.153	-0.126	-0.128	0.043			
firm appeal	-0.101	-0.173	-0.101	-0.103	0.898	0.086		
firm type B	-0.110	-0.130	-0.110	-0.112	0.985	0.880	0.041	
firm type C	-0.099	-0.115	-0.101	-0.102	0.966	0.862	0.995	0.040
Spirits								
brand market share	0.843							
brand appeal	0.723	0.836						
brand type B	0.998	0.728	0.845					
brand type C	0.997	0.730	1.000	0.849				
firm market share	-0.244	-0.176	-0.241	-0.240	0.085			
firm appeal	-0.112	-0.122	-0.117	-0.120	0.636	0.039		
firm type B	-0.255	-0.183	-0.252	-0.252	0.995	0.641	0.096	
firm type C	-0.265	-0.189	-0.264	-0.264	0.988	0.640	0.998	0.105

Table E.4: Correlations between fixed effects in the AGSU restricted sample

Notes: For brand and firm type, we use B and C to denote Bertrand and Cournot conduct, respectively. **Diagonal:** ratio of FE variances to variance of the dependent variable. **Off-diagonal:** correlation between fixed effects from regressions on samples limited to the largest connected set, brands that changed ownership, and firms with 10+ moving brands.

ically negative for spirits. Firm effects under the AGSU restrictions explain just a small part of the variance of performance measures for both beer and spirits. Therefore, the identity of the firm owning a brand explains relatively little of the variance in its market share, appeal and cost-adjusted appeal. Brand effects explain a much larger share of the overall variance. It is possible, in the presence of negative covariance between firm and brand fixed effects, for brand effects to explain more than 100% of the overall performance. We see this for beer in Table E.4.

Type of FE	# of FE	$\lambda_2$	$\Delta R^2$	Varshr	FE Corr					
Beer										
Firms (All)	464	0.000	0.005	NA	NA					
Firms (Largest connected set, AKM)	90	0.013	0.005	0.115	-0.292					
Firms (Leave-out-match, KSS)	50	0.072	0.004	0.098	-0.279					
Firms (High mobility, AGSU)	22	0.171	0.004	0.040	-0.102					
Clusters (BLM)	15	0.548	0.001	0.009	0.163					
Clusters (BLM)	10	0.748	0.001	0.009	0.135					
Clusters (BLM)	5	0.957	0.000	0.003	0.195					
Spirits										
Firms (All)	849	0.000	0.008	NA	NA					
Firms (Largest connected set, AKM)	93	0.013	0.009	0.373	-0.449					
Firms (Leave-out-match, KSS)	43	0.015	0.004	0.241	-0.365					
Firms (High mobility, AGSU)	19	0.071	0.006	0.105	-0.264					
Clusters (BLM)	15	0.348	0.001	0.029	0.074					
Clusters (BLM)	10	0.601	0.001	0.013	0.107					
Clusters (BLM)	5	0.869	0.000	0.005	0.270					

Table E.5: The explanatory power of owner fixed effects: Cournot conduct

Notes: # of FE is either number of firms or clusters.  $\lambda_2$  measures network connectivity.  $\Delta R^2$  is the difference in  $R^2$  between the full specification and one excluding firm/cluster fixed effects. Varshr is the ratio of the variance of firm/cluster FEs to the variance of brand type (ln  $\varphi_{bn}$ , conduct =Cournot). FE corr is the correlation between brand and firm/cluster FEs.

	Beer			Spirits				
Fixed effects:	b+f	b+k	bf	b+f	b+k	bf		
home	$0.475^{a}$	$0.487^{a}$	$0.485^{a}$	$0.180^{a}$	$0.175^{a}$	$0.180^{a}$		
	(0.052)	(0.049)	(0.052)	(0.044)	(0.042)	(0.044)		
distance	$-0.050^{a}$	$-0.054^{a}$	$-0.057^{a}$	$-0.064^{a}$	$-0.062^{a}$	$-0.064^{a}$		
	(0.018)	(0.016)	(0.019)	(0.015)	(0.015)	(0.015)		
common language	$0.099^{b}$	$0.103^{b}$	$0.096^{b}$	0.032	0.036	0.035		
0 0	(0.042)	(0.041)	(0.042)	(0.026)	(0.026)	(0.026)		
home (HQ)	$0.100^{c}$	0.066	0.089	$0.143^{a}$	$0.131^{a}$	$0.155^{a}$		
	(0.056)	(0.044)	(0.059)	(0.039)	(0.036)	(0.041)		
distance (HQ)	$-0.038^{b}$	$-0.027^{a}$	$-0.038^{b}$	$0.032^{a}$	$0.026^{b}$	$0.036^{a}$		
	(0.016)	(0.010)	(0.019)	(0.011)	(0.010)	(0.012)		
com. lang. (HQ)	-0.035	-0.037	-0.031	$0.053^{a}$	$0.044^{b}$	$0.053^{b}$		
	(0.036)	(0.032)	(0.038)	(0.020)	(0.019)	(0.021)		
Observations	34,675	34,675	34,675	60,624	60,624	60,624		
$\mathbb{R}^2$	0.737	0.733	0.752	0.606	0.599	0.611		
RMSE	0.208	0.208	0.203	0.210	0.211	0.208		
Notes: Standard errors in () clustered by origin-market dyads. Dependent variable: ln (obr								

Table E.6: Friction estimates, alternative heterogeneity assumptions: Cournot conduct

Notes: Standard errors in (), clustered by origin-market dyads. Dependent variable:  $\ln \varphi_{bn}$ . Market-year-product fixed effects in each regression. HQ variables determined by brand owner's headquarters country. Significance levels: 1% (*a*), 5% (*b*), and 10% (*c*).