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

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

July, 2021

Abstract

We assess the consequences for consumers in 76 countries of multinational acquisitions in beer and spirits. Outcomes depend on how changes in ownership affect markups versus efficiency. We find that owner fixed effects contribute very little to the performance of brands. On average, foreign ownership tends to raise costs and lower appeal. Using the estimated model, we simulate the consequences of counterfactual national merger regulation. The US beer price index would have been 4–7% higher without divestitures. Up to 30% savings could have been obtained in Latin America by emulating the pro-competition policies of the US and EU.

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, ERWIT, 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. Raffaele Saggio, Stephane Bonhomme, Koen Jochmans guided us on econometric methods for network data. Nathan Miller and Michele Fioretti provided helpful advice. Behrad Hafezi supplied essential research assistance.

†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] Arkolakis et al. (2018) formalize this argument as a "direct" markup effect that exceeds the more intuitive "indirect" effects coming from greater competition. 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

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

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

[3] This 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.

interactions between brands referred to as "global giants" and "local stars." The former are MNC-owned brands sold in many countries, whereas the latter are brands that obtain high market shares exclusively in their country of origin. Diageo's purchase of Yeni Raki, the most popular spirits brand in Turkey, provides a useful example. The merger raised Diageo's share of the Turkish spirits market to 63%. The pure effect of market power would not change the optimal markup for Yeni Raki if Diageo did not already have a 6% share of the Turkish market (mostly from its best-selling whisky Johnnie Walker). The combination of its global giant brand with Yeni Raki (a local star) motivates Diageo 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 the price index.

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 assume that markup determination can be adequately approximated by a Nash equilibrium, with either prices or quantities as the strategic variables.[5]

The reasoning behind our approach of estimating cost/appeal changes, but simulating

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

[5] Pinkse and Slade (2004) find that static Nash oligopoly in prices is not rejected in the British beer market. Miller et al. (2019) 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.

price changes comes from the relative strengths of our data set and our view of the most important knowledge gaps in the literature. A number of studies of mergers 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 tends to consider the US market in isolation.[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 harmed when acquisitions alter the headquarters country for their favored brands? The data we employ are uniquely well qualified for these tasks as they track brand ownership and market shares for all major markets during a decade featuring widespread ownership changes. Some of those markets start out with much higher levels of concentration than the US and are therefore more adversely impacted by mergers.

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 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

[6]The most comprehensive collection of high quality retrospective merger studies, Kwoka (2014), restricts attention to 47 studies of mergers that affected the United States.

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. The precise model we use combines elements from Atkeson and Burstein (2008), Edmond et al. (2015), Hottman et al. (2016), and Nocke and Schutz (2018b). The key features are multi-product oligopoly and nested constant elasticity of substitution (CES) demand.

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 four to six percent. Conversely, passive countries paid as much as 30% more for beer than they would have by emulating US and EU remedies. Our second contribution is to show that the owner of a brand contributes surprisingly little to its performance. Since firm effects explain just 2–7% of the variation in a brand cost-adjusted appeal, compelling a divestiture need not imply forgoing important synergies. However, a third important result is that the *geography* of ownership matters. Being owned by a firm with a faraway headquarters tends to lower cost-adjusted appeal in a market by ten to twenty percent. We believe this is the first study to estimate this negative effect of overseas ownership on the cost-adjusted appeal of a product. Finally, we show that superstar effects played little role in either beer or spirits markets over the last 12 years: Aggregate markups of the largest firms grew by putting big brands under common ownership, rather than by expanding the market shares of the high-markup brands.

In addition to the substantive findings described above, our paper makes three 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. The other method contribution is a simple way to estimate the upper level elasticity of the increasingly deployed Atkeson and Burstein (2008) model. That elasticity plays a vital role in constraining markups near monopoly. We show how to ensure that its magnitude is consistent with consolidated accounting data on markups. Third, we show how to apply recent techniques from labor economics to diagnose limited mobility bias and mitigate its im-

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

pact on the estimated contribution of firms.[8] 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 our model, and displays how oligopoly Lerner indexes vary with market share and conduct under nested CES demand. There we also 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 extra market-level variation contained in our data which 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.

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 a given period. The second set of data, created as part of this study, determines the origin of each brand. The third, also original to this study, identifies the headquarters country for the firm owning each brand. 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 (me-

[8] Jochmans 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.

dian 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.[9] 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

[9]Compustat 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.

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 quantiatively important since Compustat classifies Pernod Ricard, the second largest spirits distiller in the world, as a winery.

Table 1: Firms and their brands in the GMID beverage data

Category	**Brands**			**Firms**	**Countries**		
	All	multiple markets	multiple owners		HQ	Origin	Market
Beer	2425	368	672	464	79	93	78
Spirits	2894	598	528	849	87	106	77
Wine	1540	235	221	699	54	54	53
Water	1210	212	220	735	81	97	88
Carbonates	938	238	164	401	79	86	92
Coffee	617	153	156	390	74	79	91
Juice	1193	305	236	758	85	93	90

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. 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 different markets. The spirits category 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. For all these reasons, the rest of the paper focuses on beer and spirits, though we report regression results including other beverages in the appendix. The last three columns illustrate the diversity of headquarters countries, brand origins, and markets represented in the data.

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.[10] 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

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

Affiliations from Lexis-Nexis, and Capital IQ. Matching the name of each brand's owners 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 corporate websites, Google Images, news articles, Wikipedia, and trademark registries. For beer and spirits, the categories with the most brands, we made frequent use of crowd-sourced product rating websites.

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.[11] 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).

Diageo, the largest and most multinational of spirits makers, 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

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

Figure 1: The growth of beer multinationals

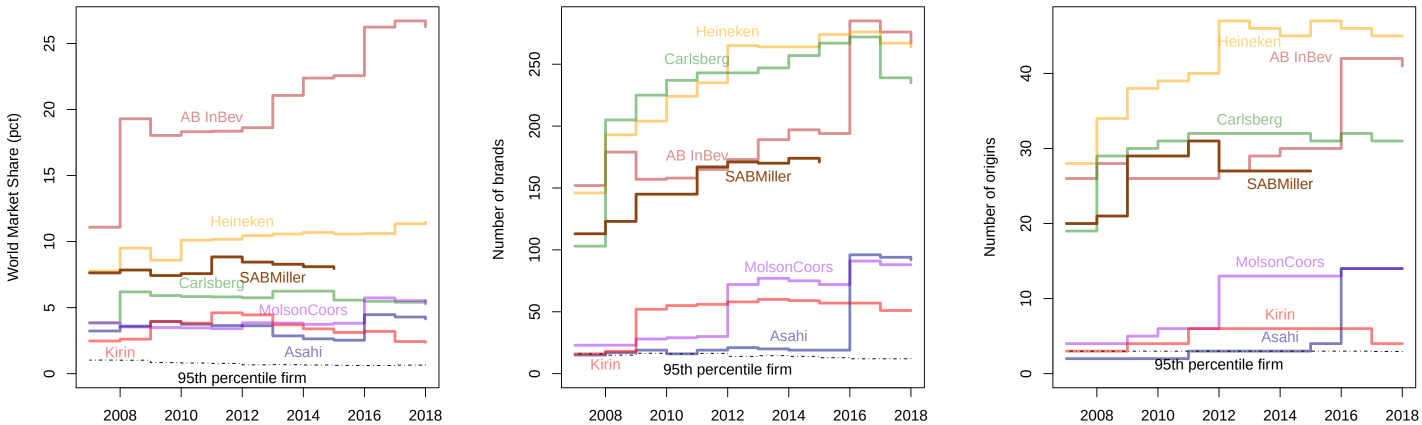

Notes: In 2008 InBev purchases Anheuser-Busch and Heineken and Carlsberg jointly purchase Scottish & Newcastle (along with BBH) and redistribute the acquired brands among themselves. In 2009 AB InBev sells off Korean and East European brands (forming Starbev) and Kirin acquires Lion (NZ). In 2012 MolsonCoors 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.

Figure 2: The growth of spirits multinationals

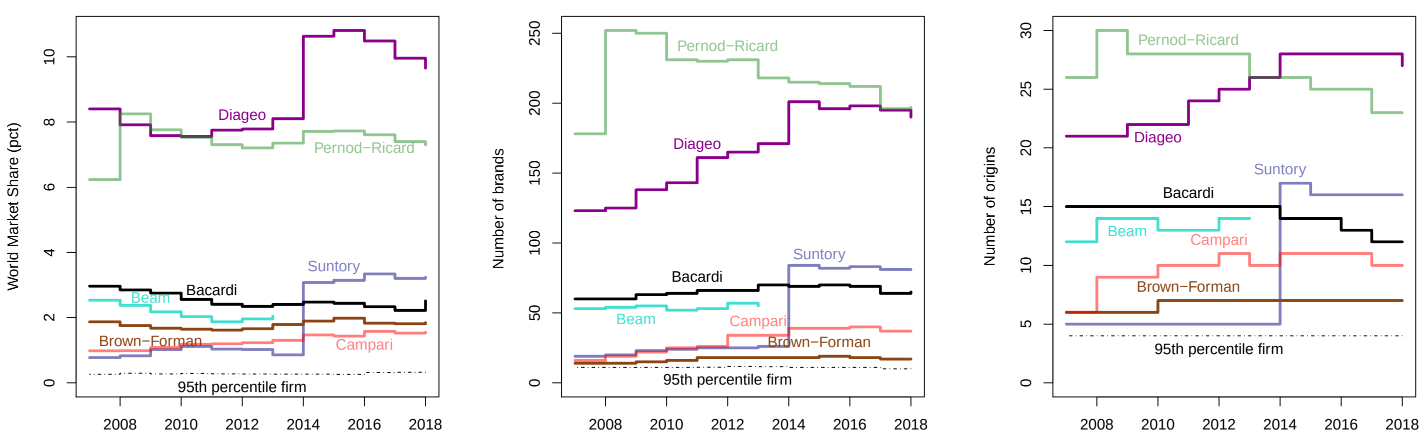

Notes: In 2008 Pernod-Ricard buys Vin & Sprit (owner of Absolut and 74 other brands). In 2014 Suntory buys Beam (which had been spun off from Fortune Brands in 2011) and Diageo buys UB Group.

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 2 displays Diageo's most prominent global giants and selects seven examples of local stars.

Table 2: A selection of Diageo brands

				Global Giants			
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			
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)

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.

The brands shown in table 2 are remarkably old, originating from 47 to 261 years ago. Not one was invented by Diageo.[12] 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.

Table 3: Statistics on global giants and local stars in 2018

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 3 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 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, deliver-

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

ing 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.[13] 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.[14] 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 exhibit a Constant Elasticity of Substitution (CES) η. Within product categories, there is a lower nest of substitution between brands with a CES of σ. This is the same preference structure as used by Atkeson and Burstein (2008), Gaubert and Itskhoki (2018), and Burstein et al. (2019), among others. Unlike those papers, we consider multiproduct firms. Adding a third nest of substitution

[13] Appendix figure A.1 visualizes these extensive margin patterns for beer, spirits and carbonates brands.

[14] The 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.

between products owned by the same firm would not alter the oligopoly markups.[15]

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 (2019b) 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 their income across a continuum of sectors, indexed $g \in [0, 1]$, with utility

$$U_n = \left[\int_0^1 Q_{gn}^{\frac{\eta-1}{\eta}} dg \right]^{\frac{\eta}{\eta-1}}, \tag{1}$$

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

$$X_{gn} = \left(P_{gn}/P_n\right)^{1-\eta} X_n \quad \text{with} \quad P_n = \left[\int_0^1 P_{gn}^{1-\eta} dg \right]^{\frac{1}{1-\eta}}, \tag{2}$$

where P_{gn} is the sectoral price index, P_n is the overall price index, and X_n is aggregate expenditure.

Inside g, the quantity index Q_{gn} is given by

$$Q_{gn} = \left[\sum_b \left(A_{bn} q_{bn}\right)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}, \tag{3}$$

where q_{bn} denotes the quantity consumed of each brand b in market n. Market-dependent brand appeal, A_{bn}, 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. In the empirics, A_{bn} is time-varying but we suppress the t subscripts in this section for simplicity. Each brand is implicitly associated with a unique sector g, so we dispense with g subscripts on all variables with b subscripts.

The market share of brand b conditional on serving market n is

$$s_{bn} = \left(p_{bn}/A_{bn}\right)^{1-\sigma} P_{gn}^{\sigma-1}, \tag{4}$$

where p_{bn} is the price of brand b in market n. Using $\mathbb{I}_{kn}$ to indicate brands offered in the

[15] See Hottman et al. (2016) and Nocke and Schutz (2018b).

market n, the sectoral price index is

$$P_{gn} = \left[\sum_k \mathbb{I}_{kn} \left(\frac{p_{kn}}{A_{kn}} \right)^{1-\sigma} \right]^{\frac{1}{1-\sigma}}. \tag{5}$$

The market share of firm f in market n, S_{fn}, sums the market shares of all the brands in the firm's portfolio ($\mathcal{F}_f$) that it offers in market n:

$$S_{fn} = \sum_{b \in \mathcal{F}_f} \mathbb{I}_{bn} s_{bn}. \tag{6}$$

As shown in tables 1 and 3, 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 $\mathbb{I}_{bnt}$. Appendix section A documents the 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 does 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 $\mathbb{I}_{bnt}$ as an exogenous characteristic of brands.

The brand-level profits earned by firm f in market n is:

$$\pi_{bn} = q_{bn}(p_{bn} - c_{bn}) = s_{bn} \frac{(p_{bn} - c_{bn})}{p_{bn}} X_{gn} = s_{bn} L_{bn} X_{gn}, \tag{7}$$

where c_{bn} is the delivered cost of a unit of brand b in market n, and $L_{bn} \equiv (p_{bn} - c_{bn})/p_{bn}$ is the Lerner index relevant in that brand-market combination. The firm maximizes the sum of π_{bn} over the set of brands it offers:

$$\Pi_{fn} = \sum_{b \in \mathcal{F}_f} \mathbb{I}_{bn} \pi_{bn}, \tag{8}$$

3.2 Markups for different conduct assumptions

The pricing strategy of firms conforms with the "small in the large but large in the small" assumption of Atkeson and Burstein (2008) and Neary (2016). Firms realize and account for their influence on the price index within a sector (large in the small), but treat the aggregate expenditure and price levels (X_n and P_n) as given (small in the large).

We find it useful to express price-cost relationships in two different ways, both of which we refer to as "markups." To see how costs affect prices and how markups affect market shares, it is useful to work with $\mu \equiv p/c$, the price/cost markup. When computing profits on the other hand, the Lerner index is more convenient as seen in equation (7). The first order conditions for maximization of equation (8) yield equations for the brand-level price/cost markup and the Lerner index expressed as functions of the *firm-level* perceived elasticity of demand, ϵ_{fn}:

$$\mu_{bn} = \mu_{fn} = \frac{\epsilon_{fn}}{\epsilon_{fn} - 1}, \quad \text{and} \quad L_{bn} = L_{fn} = \frac{1}{\epsilon_{fn}}, \quad \forall b \in \mathcal{F}_f. \tag{9}$$

Prices can be expressed in terms of either markup:

$$p_{bn} = \mu_{fn} c_{bn} = c_{bn}/(1 - L_{fn}). \tag{10}$$

The property that, under CES demand, firms equate markups across all their products was derived by Feenstra (2003, p. 267) and features prominently in Hottman et al. (2016) and Nocke and Schutz (2018b).[16]

The functional form of markups depends on the assumed mode of oligopoly conduct. The Lerner indices implied by the two standard conduct assumptions are

$$\underbrace{L_{fn} = \frac{1}{\sigma - (\sigma - \eta) S_{fn}}}_{\text{Bertrand}} \quad \text{and} \quad \underbrace{L_{fn} = \frac{1}{\sigma} - \left(\frac{1}{\sigma} - \frac{1}{\eta}\right) S_{fn}}_{\text{Cournot}}. \tag{11}$$

For the set of brands that belong to the monopolistically competitive fringe, denoted with subscript 0, we have the usual constant-markup rule, with $L_{0n} = 1/\sigma$.

A major attraction of the CES oligopoly model is that it provides simple expressions for markups that rely on observable firm-level market shares, to be combined with two parameters, σ and η. We now describe how we obtain those two critical elasticities.

3.3 Matching elasticities to moments

Industrial organization economists have already devoted considerable efforts to the estimation of brand-level own-price elasticities for the very products we study. We treat those estimated elasticities as moments to pin down σ_g for each of the categories we consider.

The underlying papers summarized in table 4 report the mean or median of brand-level own-price elasticities, ϵ_b, estimated from the demand side of their models before

[16] It contrasts sharply with the case of multi-product firms facing linear demand, as analyzed by Mayer et al. (2014).

Table 4: Estimates of own-price elasticities and implied elasticities of substitution

Product group	Mean σ_g	Mean own elas. (ϵ_b)	# Estimates	# Papers
Beer	4.49	4.48	9	5
Spirits	3.38	3.37	9	2

Sources: For beer, Asker (2016), De Loecker and Scott (2016), Hausman et al. (1994), Miller and Weinberg (2017), Pinkse and Slade (2004). For spirits, Miravete et al. (2018) and Conlon and Rao (2015).

imposing a specific conduct assumption. Those demand elasticities cannot be interpreted as direct estimates of the elasticity of substitution σ_g because of non-negligible market shares. Instead, we invert the brand-level formula for CES own-price elasticity $\epsilon_b = \sigma_g - (\sigma_g - \eta)s_b$ and solve for σ_g as a function of a moment, $m_g()$ (either mean or the median, depending on the paper), of the estimated demand elasticities in the category:

$$\sigma_g = \frac{m_g(\epsilon_b) - m_g(s_b)\eta}{1 - m_g(s_b)}.$$

Due to constraints imposed by the existing empirical literature, the σ_g of 4.5 for beer and 3.4 for spirits are assumed to be constant over time and across countries.

In contrast to the abundance of high quality brand-level elasticity estimates, the literature does not provide obvious candidates for η, the CES between product categories. Atkeson and Burstein (2008), the pioneering work using nested CES oligopoly, impose $\eta = 1.01$ and consider $\eta = 1.5$ in a sensitivity analysis. Burstein et al. (2019) exploit a linear relationship between the inverse of the harmonic mean markup and the Herfindahl index to estimate a parameter corresponding to $\frac{1}{\sigma} - \frac{1}{\eta}$ using cross industry variation. They impose $\sigma = 7$ and this leads to an η estimate of 1.7. Using $\sigma = 4.5$ (the value for beer) would imply $\eta = 1.5$. Because η is so important in our quantification of markups, our η estimate should conform with markup data from the industries we focus on, beer and spirits.

We calibrate η to provide the best fit between theoretical and accounting markups. If there are constant returns to scale and no fixed costs, then the profit to sales ratio can be expressed as $(pq - cq)/pq = (p - c)/p = L$. Accounting data are generally unavailable at the market level because firms report their "consolidated" accounts, aggregating over all markets they serve. Therefore, our accounting measure of the firm-level Lerner index, denoted L_f^A, is the ratio of a firm's worldwide profits over worldwide sales. The theoretical counterpart to L_f^A, denoted L_f, must therefore also be constructed by aggregating profits implied by the model in each country. Since profit in a market is given by $L_{fn}S_{fn}X_{gn}$, the aggregate theoretical markup is just a sales-weighted average of the firm's theoretical

Lerner indexes in each country:

$$L_f = \sum_n \omega_{fn} L_{fn}, \quad \text{where} \quad \omega_{fn} = \frac{S_{fn} X_{gn}}{\sum_n S_{fn} X_{gn}}. \tag{12}$$

To calculate the accounting markup, L_f^A, in a way that corresponds to the theoretical markup, we need to purge accounting measures of costs from their fixed cost components. However, as discussed in Syverson (2019b), accounting expense categories do not map cleanly to economic concepts of fixed and variable costs. Most firms report two major categories of operating expenses: "cost of goods sold" (COGS) and "selling, general, and administrative" (SGA) expenses.[17] The accounting markup expressed in terms of the underlying Compustat variables is

$$L_f^A = \frac{\text{sale}_f - \vartheta_1 \text{cogs}_f - \vartheta_2 \text{xsga}_f}{\text{sale}_f},$$

where ϑ_1 and ϑ_2 denote the fractions of each cost category assumed to be marginal costs. As in De Loecker et al. (2020), we take COGS to be entirely variable costs, implying $\vartheta_1 = 1$. For ϑ_2 we consider two bounding cases. Our conservative markup measure treats all of SGA as variable costs ($\vartheta_2 = 1$), leading to our lower bound on accounting markups. Since SGA includes cost categories such as administration and R&D that seem like classic examples of overhead costs, the conservative markups are likely too low.[18] On the other hand, SGA includes distribution costs, which almost certainly vary with the amount of beer being distributed. AB InBev's annual reports provide a distinct line for distribution costs. On average, they comprise 32% of SGA from 2008 to 2018. Hence, we calculate a liberal markup deducting only $\vartheta_2 = 0.32$ of SGA.

The ω_{fn} in equation (12) are data. The L_{fn} markup formulae in equation (11) use the S_{fn} data, the calibrated σ_g (now taken as known), leaving a single unknown parameter, η. The loss function used to calibrate η is the squared deviations between theory and accounting markups:

$$\ell(\eta) = \sum_f \sum_t \left(\left[\frac{L_{ft}^{\text{Bertrand}}(\eta) + L_{ft}^{\text{Cournot}}(\eta)}{2} \right] - L_{ft}^A \right)^2. \tag{13}$$

[17] In the few instances in Compustat where xsga is incomplete, we replace it with operating expenses (xopr) minus cogs.

[18] Administrative expenses constitute a small share of SGA for the four companies that report them separately. Their share of SGA over 2008–2018 are 20% for Carlsberg and AB InBev, 14% for Royal Unibrew and 21% for Tsingtao.

We compute L^A_{ft} and the conduct-specific L_{ft} for the 14 largest publicly traded multinationals in beer and spirits (shown in Figures 3) over the 2007–2018 period. There are 157 observations (some firms are absorbed via mergers, leading to an unbalanced panel). For the estimation of η we set $\vartheta_2 = 0.66$, the average of the conservative and liberal values. The η that minimizes equation (13) is 1.62, which corresponds to a monopoly Lerner index of 62%.

Figure 3: Oligopoly markups for Bertrand and Cournot, compared to accounting data

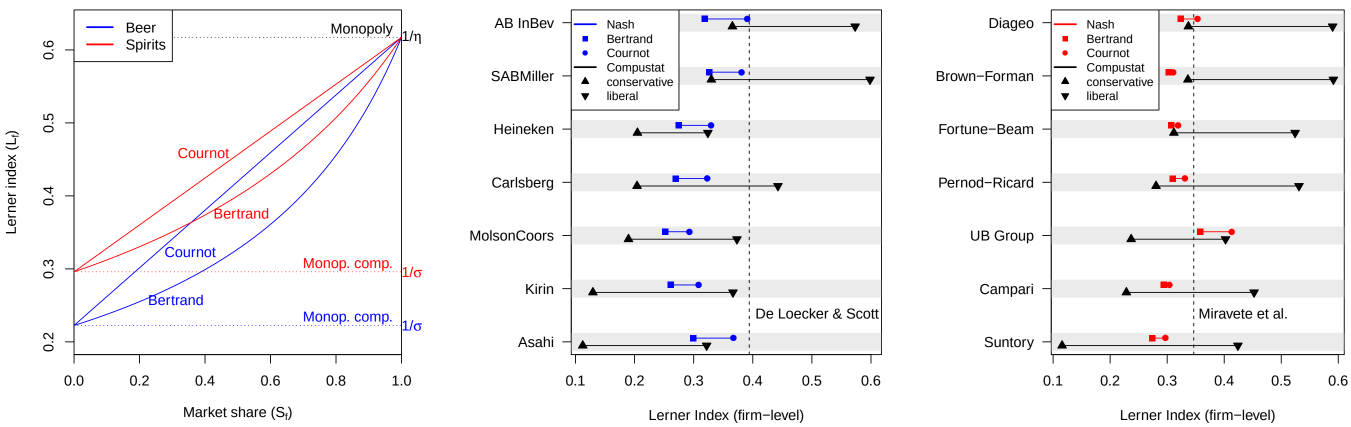

(a) Theory (b) Beer data (c) Spirits data

With σ_g and η in hand, we can graphically compare the theoretical markups to those obtained from accounting data. Figure 3(a) plots the Lerner index functions under Bertrand and Cournot conduct assumptions. The blue lines use our estimate for beer ($\sigma = 4.5$) whereas the red line uses our estimate for spirits ($\sigma = 3.4$). In both Cournot and Bertrand, L_f ranges from $1/\sigma$ for $S_f = 0$ (the monopolistic competition benchmark) to $1/\eta = 0.62$ for $S_f = 1$ (monopoly). For a given product, the Lerner index for Bertrand lies under the corresponding index for Cournot for $0 < S_f < 1$.

Figures 3(b) and 3(c) display for 2013 (before several large mergers) the Bertrand to Cournot range of Lerner indexes (in blue for beer and red for spirits). Below each theoretical interval, we show the range between our conservative and liberal bounds for accounting markups (L^A_f, in black). As a third type of comparison, vertical dashed lines display the average markups reported by De Loecker and Scott (2016) for beer and Miravete et al. (2018) for spirits. Both papers use random-coefficients logit demand models and De Loecker and Scott (2016) also provides estimates based on the De Loecker and Warzynski (2012) method.[19]

[19]Miravete et al. (2018) report weighted average Lerner indexes obtained through the standard IO demand-side approach. De Loecker and Scott (2016) report sales-weighted price-cost markups (μ) rang-

There are three salient points in the markup figures. The accounting and theory intervals overlap for every beer maker and for all but one spirits maker (Brown-Forman). The theoretical markups (based on calibrated σ_g and η) are broadly consistent with the accounting data, showing that the CES oligopoly model passes a first stress test of its suitability for the two industries we consider. The second point is that markups in the nested CES model are reasonably close to those obtained using methods preferred in the IO literature. The beer estimates of De Loecker and Scott (2016) are on the high side but they are sales-weighted and apply to the highly concentrated US market. The third noteworthy aspect of the figure is that Bertrand and Cournot theoretical markups differ less from each other than the reasonable range for accounting markups. Neither conduct assumption can be ruled out, so we will consider results for both.

3.4 Backing out cost-adjusted appeal (brand type)

Borrowing from Nocke and Schutz (2018b), 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 will also refer to as cost-adjusted appeal.[20] With estimates of the demand elasticities, data on brand sales shares in a market allow us to back out all the φ_{bn} up to a normalization. The n subscripts are important here because the data reveal large variation in φ_{bn} across markets.

Substituting for equilibrium price and then inverting equation (4) we obtain

$$\varphi_{bn} = s_{bn}^{1/(\sigma-1)} \mu_{fn} P_{gn}. \tag{14}$$

In order to isolate brand type as a function of observables, we need to eliminate P_{gn}. This can be accomplished by dividing by the φ of any other brand facing the same price index. We divide by the geometric mean of the individually listed brands in a market so that our inferred brand types will have the same normalization as Hottman et al. (2016) use for inferring brand appeal. Following Hottman et al. (2016), a tilde over a variable denotes its geometric mean over the relevant market-year (specified in its subscript). By the properties of geometric means, $\tilde{\varphi}_{gnt} = \tilde{s}_{gnt}^{1/(\sigma-1)} \tilde{\mu}_{gnt} P_{gnt}$. Incorporating t into (14) and

ing from 1.6 to 1.7 in different specifications of the demand-side method and 1.65 using the production function approach. We transform the average μ to Lerner equivalents by $L = 1 - 1/1.65 = 0.39$.

[20] Melitz (2003) points out the isomorphism in a model of CES single-variety monopolistic competition. Nocke and Schutz (2018b) 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.

dividing by the geometric mean,

$$\check{\varphi}_{bnt} = \frac{\varphi_{bnt}}{\tilde{\varphi}_{gnt}} = \left(\frac{s_{bnt}}{\tilde{s}_{gnt}}\right)^{1/(\sigma-1)} \frac{\mu_{fnt}}{\tilde{\mu}_{gnt}}. \tag{15}$$

Markups for all listed brands are functions of the demand parameters and firm-level market shares and therefore can obtained by applying a conduction assumption inside equation (11), and using $\mu_{fn} = 1/(1 - L_{fn})$.

Following Khandelwal et al. (2013) and Redding and Weinstein (2018), one can infer the relative demand shifter from the data and an estimate of σ:

$$\check{A}_{bnt} = \frac{A_{bnt}}{\tilde{A}_{gnt}} = \left(\frac{s_{bnt}}{\tilde{s}_{gnt}}\right)^{1/(\sigma-1)} \frac{p_{bnt}}{\tilde{p}_{gnt}}. \tag{16}$$

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 Estimation of ownership effects on brand performance

The focus in this section is to estimate the impact of firm ownership on brand performance (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.

4.1 Estimating equations

We now derive from the model the equations we estimate. There are three mappings that we use repeatedly in the specifications:

- $o(b, t)$ maps a brands to its *owner* in year t.[21]

[21]There are some brands, e.g. Fosters, whose owner varies across countries. We omit the n subscript from

- $h(f)$ maps a firm to location of its *headquarters*.
- $i(b)$ maps a brand to its *origin*, the country where the brand was introduced.

Substituting for price in equation (4), applying the definition of brand type, and taking logs, we have

$$\ln s_{bnt} = (\sigma - 1)\left[\ln \varphi_{bnt} - \ln \mu_{o(b,t)nt}\right] + (\sigma - 1)\ln P_{gnt}. \tag{17}$$

The last term in this equation can be eliminated with fixed effects defined at the product-market-year level. The delivered cost-adjusted appeal, φ_{bnt} can be further decomposed into a brand-specific term, φ_b^B, an owner-specific term, $\varphi_{o(b,t)}^F$, a friction between brand origin and market denoted $\delta_{i(b)n}^B$, a friction between the current brand owner's headquarters and market denoted $\delta_{h(o(b,t))n}^F$, and a residual.

$$\ln \varphi_{bnt} = \ln \varphi_b^B + \ln \varphi_{o(b,t)}^F + \ln \delta_{i(b)n}^B + \ln \delta_{h(o(b,t))n}^F + \varepsilon_{bnt}. \tag{18}$$

The δ^B and δ^F capture the impact of observable frictions on φ_{bnt}. The δ 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 $\delta_{i(b)n}^B$ and $\delta_{h(o(b,t))n}^F$. The first, $\text{home}_{i(b)n}$, takes a value of 1 for brands sold in their country of origin ($i = n$). The second, $\text{home}_{h(o(b,t))n}$, 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.

We can now be more concrete about the contents of the residual ε_{bnt}. All shocks to appeal or costs that are specific to the brand-market dyad enter there. In addition, it includes all the *unobserved* determinants of the δ frictions. Moreover, ε_{bnt} captures cost determinants related to the location of production—which our data does not report. The simplest case to consider are brands of Scotch Whisky or Champagne that by law must be produced in origin country i. In such cases the coefficient on log distance captures not only the elasticity of appeal with respect to distance, but also the elasticity of iceberg transport costs (from Scotland or France to market n). More generally, the estimates on each friction determinant will be increasing in multinational production costs associated with serving remote markets (either by horizontal investment or export platforms). Such effects would be most likely to show up in the hn dimension if management of overseas production is based on the brand owner's headquarters.

$o(b, t, n)$ in the notation, but take it into account in the estimation and counterfactuals.

The final estimating equation for cost-adjusted appeal uses our inferred values, $\breve{\varphi}_{bnt}$ from (15) in place of the unobservable φ_{bnt}.

$$\ln \breve{\varphi}_{bnt} = \mathrm{VFE}^B_b + \mathrm{VFE}^F_{o(b,t)} + \mathrm{VFE}_{gnt} + \mathbf{X}'_{i(b)n}\mathbf{d}^B + \mathbf{X}'_{h(o(b,t))n}\mathbf{d}^F + \varepsilon_{bnt}, \tag{19}$$

where $\mathbf{X}$ comprises home, distance, and common language, measured with respect to the brand origin when subscripted with i and with respect to HQ when subscripted with h. The VFEs (varphi fixed effects) have structural interpretations as $\ln \varphi^B_b$, $\ln \varphi^F_{o(b,t)}$, and $-\ln \tilde{s}_{gnt}/(\sigma-1) - \ln \tilde{\mu}_{gnt}$. To determine the effect of each friction variable working through the demand side alone, we also estimate a version of equation (19) where $\ln \breve{A}_{bnt}$ replaces $\ln \breve{\varphi}_{bnt}$ as the dependent variable. The differences between the coefficients in those two regressions correspond to the cost effects of each friction determinant.

The key identifying assumption for the estimating equation (19) is that the expectation of ε_{bnt} is zero, *conditional* on the firm and brand fixed effects and the frictions. One threat to this assumption would be interactions between unobserved brand and firm characteristics. While our baseline specification assumes that any such interactions are orthogonal to the friction determinants, we also consider a specification that allows for a general pattern of firm-brand interactions.

The primitive determinant of brand market shares in equation (17) is the brand's cost-adjusted appeal within the market, φ_{bnt}. It is also interesting to estimate the impact of frictions on the other variable featured in the same equation, the markup. We therefore regress log markups on the same set of fixed effects and frictions, yielding

$$\ln \mu_{bnt} = \mathrm{MFE}^B_b + \mathrm{MFE}^F_{o(b,t)} + \mathrm{MFE}_{gnt} + \mathbf{X}'_{i(b)n}\mathbf{g}^B + \mathbf{X}'_{h(o(b,t))n}\mathbf{g}^F + \upsilon_{bnt}. \tag{20}$$

In this regression, the coefficients do not reveal structural parameters because of the non-linear mapping from frictions to market shares and from market shares to markups. The markup fixed effects (MFE) also do not map in any simple way to structural parameters.

Substituting the cost-adjusted appeal and markup equations into (17), we have the estimable log market share equation:

$$\ln s_{bnt} = \mathrm{SFE}^B_b + \mathrm{SFE}^F_{o(b,t)} + \mathrm{SFE}_{gnt} + \mathbf{X}'_{i(b)n}\mathbf{r}^B + \mathbf{X}'_{h(o(b,t))n}\mathbf{r}^F + \xi_{bnt}. \tag{21}$$

The additive-in-logs structure implies that market share friction coefficients are algebraically tied to the $\ln \breve{\varphi}_{bnt}$ and $\ln \mu_{bnt}$ coefficients via $\mathbf{r} = (\sigma - 1)(\mathbf{d} - \mathbf{g})$. Similarly, the coefficients on $\ln \breve{\varphi}_{bnt}$ and $\ln \mu_{bnt}$ for different conduct assumptions are linked through equation (15): the difference between friction coefficients on the Cournot and Bertrand versions of φ is constrained to equal the corresponding difference in μ coefficients. The error term for market shares relates back to the two previous error terms via $\xi_{bnt} = (\sigma - 1)(\varepsilon_{bnt} - \upsilon_{bnt})$. Thus, this error captures brand-market idiosyncratic shocks (to appeal and cost), unobserved friction determinants, and specification error in the markup equation.

4.2 Baseline estimation results

Table 5 reports results for regressions that pool beer and spirits brands. The most striking result is the huge advantages held by home-origin brands. Since $\exp(1.029) \approx 2.8$, home increases market share by 180%. Our estimate of the home advantage for beer and spirits brands is somewhat larger than the 126% estimate for car brands obtained in Head and Mayer (2019a). Distance from brand origin also reduces market share, with an elasticity of -0.12. Head and Mayer (2019a) estimate a 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 $\exp(0.353) - 1$ = 42% 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 25% price reduction.[22]

The R^2 of the brand type estimations in table 5 are 0.6 (for both conduct assumptions), indicating that idiosyncratic shocks explain about 40% of the variation in $\ln \breve{\varphi}_{bnt}$. This finding motivates the usefulness of exact hat algebra for counterfactuals since EHA implicitly takes into account the unobserved determinants of market share that are invariant to the counterfactual.

The pooled regressions in Table 5 estimate the effect of frictions averaging over 12 years and two products. To assess how beer and spirits home bias compare to each other, and how they evolve over time, we estimate a model for each product 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 (φ).[23] 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

[22]Goldberg 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.

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

Table 5: Brand performance regressions: Beer and Spirits

			Bertrand		Cournot	
	$\ln s_{bn}$	$\ln A_{bn}$	$\ln \varphi_{bn}$	$\ln \mu_{bn}$	$\ln \varphi_{bn}$	$\ln \mu_{bn}$
home	1.029^a	0.223^a	0.353^a	0.016^a	0.367^a	0.030^a
	(0.135)	(0.072)	(0.051)	(0.004)	(0.053)	(0.006)
distance	-0.121^a	0.036	-0.044^a	-0.002^c	-0.046^a	-0.004^b
	(0.037)	(0.022)	(0.015)	(0.001)	(0.015)	(0.002)
common language	0.047	−0.054	0.008	0.0001	0.008	0.0003
	(0.078)	(0.050)	(0.031)	(0.002)	(0.032)	(0.003)
home (HQ)	0.354^a	0.104^c	0.179^a	0.031^a	0.204^a	0.056^a
	(0.106)	(0.061)	(0.042)	(0.003)	(0.043)	(0.006)
distance (HQ)	0.019	0.009	0.013	0.001	0.012	−0.001
	(0.033)	(0.020)	(0.013)	(0.001)	(0.014)	(0.001)
com. lang. (HQ)	0.114^c	0.048	0.052^b	0.003	0.056^b	0.006
	(0.062)	(0.038)	(0.025)	(0.003)	(0.026)	(0.004)
Observations	95,299	95,299	95,299	95,299	95,299	95,299
R^2	0.657	0.653	0.596	0.900	0.603	0.859

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).

effects do not depend on conduct, since they are extracted directly as demand shifters.

As seen in panel (a) of Figure 4, the total effect of being a home origin beer brand is equivalent to a 55–60% 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 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.[24]

[24]This explanation is further supported by Tables 7 and C.8, where the distance coefficients for beer are more than twice as large as those for spirits.

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

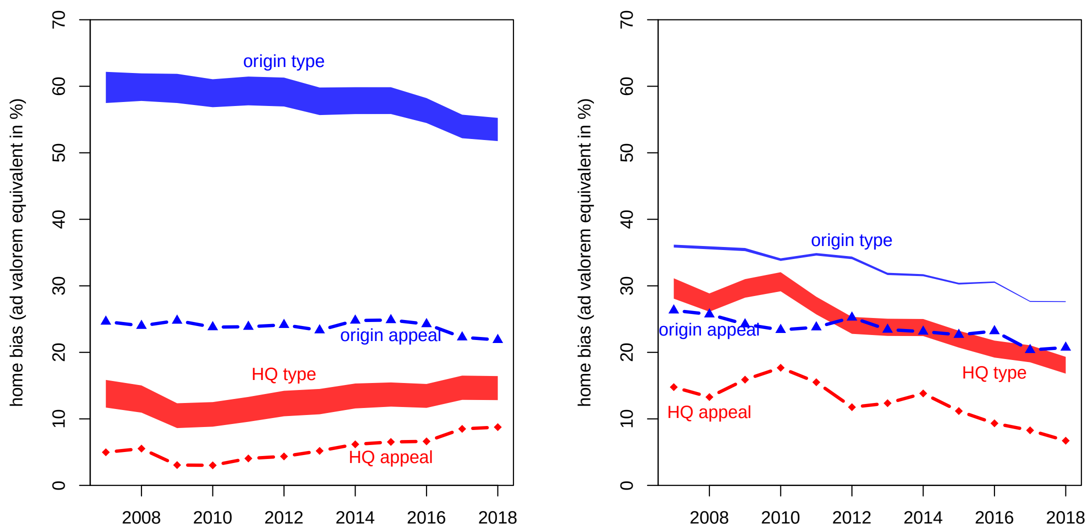

(a) Beer (b) Spirits

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

HQ-related home advantage is estimated as equivalent to around a 10–15% tariff for beer, and 20%–30% in spirits. 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 VFE_f in equation (19) is $\ln \varphi^F_{o(b,t)}$. A transfer of b to a new owner in period $t+1$, raises cost-adjusted appeal by $\ln \varphi^F_{o(b,t+1)} - \ln \varphi^F_{o(b,t)}$. 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.

4.3 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 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 5.[25] 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 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 λ_2, is calculated as the smallest non-zero eigenvalue of the normalized weighted Laplacian of the induced network.[26] 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 B.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

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

[26]Appendix section B provides greater detail on this procedure.

(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 brands the firm owns in the first period (2007 for most firms).[27]

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.[28] The second step of KSS constructs a finite sample unbiased variance estimator that is computed by repeatedly leaving a single brand-market-year observation out of the sample.

Table 6 summarizes our results on the firm effects for beer and spirits.[29] The incremental R^2 for firm fixed effects is just 0.007 for both beverages. That is, firms add very little explanatory power to a specification that already includes brand effects and the six

[27] As 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).

[28] In the network illustrated in figure B.1, Fosters is a bottleneck brand.

[29] Results for Cournot conduct are very similar to the Bertrand results shown here, so they are relegated to Appendix table C.7.

Table 6: The explanatory power of owner fixed effects

Type of FE	# of FE	λ_2	ΔR^2	Varshr	FE Corr
		Beer			
Firms (All)	464	0.000	0.007	NA	NA
Firms (Largest connected set, AKM)	90	0.013	0.008	0.359	-0.497
Firms (Leave-one-out, KSS)	49	0.071	0.003	0.059	-0.129
Firms (High mobility, AGSU)	22	0.171	0.004	0.036	-0.069
Clusters (BLM)	15	0.461	0.001	0.026	0.110
Clusters (BLM)	10	0.548	0.001	0.029	0.183
Clusters (BLM)	5	0.618	0.001	0.024	0.204
		Spirits			
Firms (All)	849	0.000	0.007	NA	NA
Firms (Largest connected set, AKM)	93	0.013	0.007	0.231	-0.500
Firms (Leave-one-out, KSS)	41	0.010	0.013	0.098	-0.146
Firms (High mobility, AGSU)	18	0.071	0.006	0.065	-0.035
Clusters (BLM)	15	0.426	0.002	0.051	0.155
Clusters (BLM)	10	0.436	0.002	0.054	0.169
Clusters (BLM)	5	0.904	0.001	0.022	0.292

Notes: # of FE is either number of firms or clusters. λ_2 measures network connectivity. ΔR^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 ($\ln \breve{\varphi}_{bn}$, conduct =Bertrand). FE corr is the correlation between brand and firm/cluster FEs. References for AKM, AGSU, BLM, KSS given in text.

friction variables. We now turn to the standard way of measuring firm value added since AKM: the variance of the firm fixed effects divided by the variance of the dependent variable.

The largest connected set includes 20% of the firms in beer and 11% in spirits. However, these firms account for the majority of world sales.[30] The second row for each beverage gives a startling—but misleading—impression of the importance of firms and it suggests strongly negative assortative matching between owners and brands. The λ_2 connectivity of both sets is just 0.01, compared to $\lambda_2 = 1.00$ for a fully connected network.[31] Moving to the KSS leave-one-out estimator, the number of firms falls to 49 (beer) or 41 (spirits). As expected, based on the results of Kline et al. (2020), the variance share of firm fixed effects falls sharply, as does the estimated amount of negative assortative matching.[32]

The subsequent rows of Table 6 establish that when λ_2 connectivity exceeds 0.07, the variance share of firm fixed effects shrinks to the 0.02–0.07 range for both products. Moreover, the strong negative assortative matching is revealed to be an artifact of low connectivity. Restricting to the set of high mobility firms, raises λ_2 to 0.17 for beer and 0.07 for spirits, which is sufficient to put the variance share below 4% and 7%, for beer and spirits, respectively.[33] The firm effects estimated in this sample have negligible correlations with their corresponding brand effects.[34]

The group fixed effects method eliminates the suspicious negative assortative matching. The results shown in the Clusters (BLM) rows of Table 6 convey a common message about firm value-added whether we use $K = 10$ as in BLM, $K = 5$, or $K = 15$. In each case, connectivity is over 0.4 and the value added of owners is 2–5% of the variance in brand type. Although group effects work by reducing dimensionality, they still capture a substantial majority of the *between firm* variance in brand type. With $K = 10$, the clusters account for 62% of the firm-mean variance in $\ln \breve{\varphi}_{bn}$ for beer and 57% for spirits. Raising K to 15 makes little difference.

Table 7 shows how estimates of the frictions change as we deviate from the baseline specification of additively separable brand and firm effects (a la AKM). Columns (1) and (4) show, separately for beer and spirits brands, the baseline specification. Columns (2) and (5) show the clustered (or group) fixed effects. This is the same regression as the one

[30]Table B.1 shows that sales of the largest component accounts for 80% of beer and 58% of spirits.

[31]Interestingly, 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).

[32]We implement the KSS estimator using the option to restrict the sample to moving brands.

[33]The remaining 21 individual beer makers still account for a respectable 71% of total beer sales, while the remaining 17 spirits makers account for 42% of total spirits sales.

[34]The device of the superfirm plays a quantitatively important role, especially for spirits.

Table 7: Brand type regressions with alternative heterogeneity assumptions

	Beer			Spirits		
Fixed effects:	$b+f$	$b+k$	bf	$b+f$	$b+k$	bf
home	0.444^a	0.465^a	0.451^a	0.279^a	0.270^a	0.277^a
	(0.054)	(0.053)	(0.055)	(0.067)	(0.065)	(0.068)
distance	-0.073^a	-0.063^a	-0.081^a	-0.032^c	-0.031^c	-0.032^c
	(0.018)	(0.017)	(0.019)	(0.019)	(0.018)	(0.019)
common language	0.091^b	0.104^a	0.086^b	-0.019	-0.017	-0.020
	(0.041)	(0.039)	(0.041)	(0.039)	(0.038)	(0.040)
home (HQ)	0.103^c	0.060	0.096^c	0.210^a	0.201^a	0.226^a
	(0.053)	(0.044)	(0.056)	(0.056)	(0.052)	(0.059)
distance (HQ)	-0.032^c	-0.033^a	-0.030	0.029^c	0.028^c	0.030^c
	(0.016)	(0.012)	(0.020)	(0.017)	(0.015)	(0.017)
com. lang. (HQ)	-0.026	-0.035	-0.014	0.075^b	0.067^b	0.075^b
	(0.036)	(0.032)	(0.039)	(0.030)	(0.029)	(0.031)
Observations	34,675	34,675	34,675	60,624	60,624	60,624
R^2	0.736	0.730	0.748	0.549	0.544	0.553
RMSE	0.236	0.237	0.232	0.385	0.384	0.382

Standard errors in (), clustered by origin-market dyads. Dependent variable: $\ln \check{\varphi}_{bn}$. Market-year-product fixed effects in each regression. HQ variables determined by brand owner's head-quarters country. In the second and fifth columns, k corresponds to the group FE ($K = 10$). Significance levels: 1% (a), 5% (b), and 10% (c).

reported for 10 clusters in Table 6. Also, this specification provides the friction and group fixed effect estimates underlying Figure 5 and the counterfactual exercises.

Columns (3) and (6) of table 7 show a new specification that replaces the additive b and f fixed effects with interactive bf fixed effects. If firm-brand "match effects" 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. Analogously to the approach taken by Card et al. (2013), we respond to this concern by estimating a specification with a full set of brand-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 cross-market and cross-time variation. Since this specification nests the $b + f$ specification, the R^2 necessarily rises. However, the change is very small (0.012 for beer, 0.004 for spirits) and the root mean squared error (RMSE) hardly declines. The implied standard deviation of the match effect is just 0.043 for beer and 0.047 for spirits.[35] The friction estimates themselves change very little across the three specifications, suggesting that the orthogonality assumption for the match effects is not strongly violated. We reproduce this set of regressions as Table C.8 in the Appendix, with Cournot $\breve{\varphi}_{bn}$ as the dependent variable. The friction coefficients are slightly larger and more statistically significant under Cournot, but the pattern of changes in R^2 and RMSE are essentially the same.

Figure 5 visualizes the distributions of changes in $\breve{\varphi}_{bn}$ that our estimates imply to have occurred as a consequence of the observed set of brand ownership changes. The blue densities shows changes in $\breve{\varphi}_{bn}$ 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 cost-adjusted appeal when domestic firms purchase foreign-owned brands—is rare.

The red densities in Figure 5 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 and Grupo Modelo (Corona). The difference in $\ln \breve{\varphi}_{bn}$ between the groups to which AB InBev and Anheuser Busch respectively belong corresponds to a 0.02 log point reduction of Budweiser's brand type. On the other hand, when Heineken

[35] As in Card et al. (2013), we calculate this as the square root of the difference between the squared RMSEs of the bf and $b + f$ columns.

Figure 5: How ownership changes affect brand type (φ_{bn})

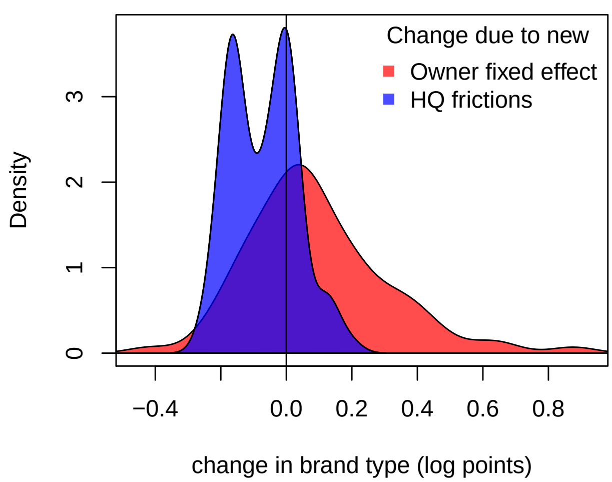

(a) Beer firms (LCS)

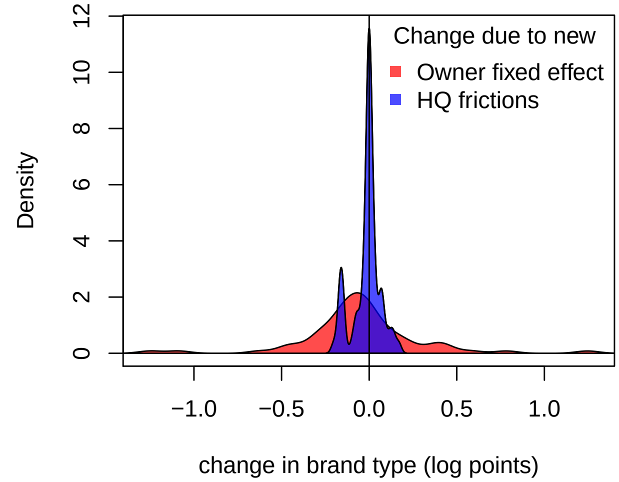

(b) Spirits firms (LCS)

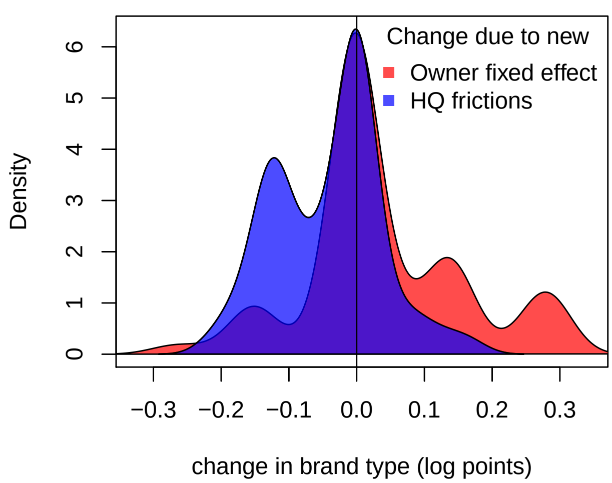

(c) Beer firm clusters ($K = 10$)

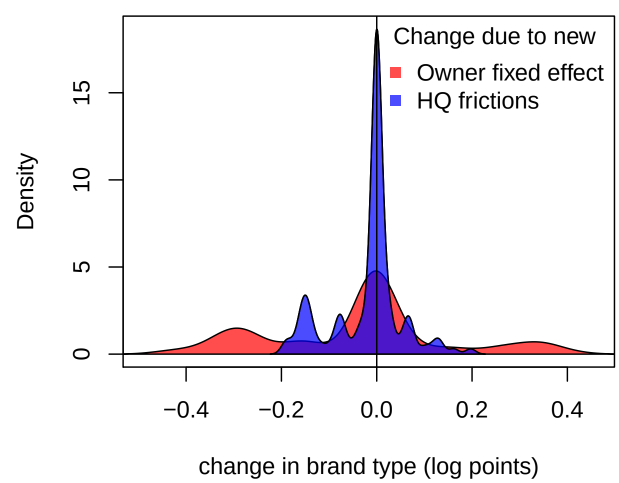

(d) Spirits firm clusters ($K = 10$)

bought Lagunitas, the latter benefited from a a 0.28 improvement in $\ln \check{\varphi}_{bn}$. The Belgian craft brewery Bosteels made the same large move when AB InBev acquired it in 2016.[36] Another important finding displayed in figure 5 is that the range of group effects is about 0.7 for firm-clusters in beer which is much smaller than the 1.3 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. Since synergies of this form would not influence brand market shares, they would not influence the price outcomes of ownership changes. Hence, we do not need to take a stance on them in the counterfactuals when considering the consequences of mergers on the consumer surplus, the exercise to which we now turn.

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.

In addition to taking into account how alternative ownership patterns affect firm level

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

market shares and hence their optimal markups, we also account for the changes in brand type (φ_{bn}) implied by the counterfactual ownership, using estimates from columns 2 and 5 (beer and spirits, respectively) of Tables 7 (Bertrand) or C.8 (Cournot), as illustrated in Figure 5(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 set of brand portfolios for each firm which we denote as $\mathcal{F}'_f$. 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. So far as we know, this is the first application extending EHA to incorporate oligopoly markup adjustment, which permits counterfactual merger analysis. With EHA, only changes in φ_{bn} need to be specified and they are obtained from the regressions of the previous section. We denote (proportional) changes for all variables with hat notation, for instance $\hat{s}_{bn} \equiv s'_{bn}/s_{bn}$, where s'_{bn} is the new level of s_{bn} under the counterfactual change.

The ownership changes that we simulate imply a change in the firm-destination market shares, for two reasons: 1) the changes in the number and identity of brands owned, 2) the changes in equilibrium market shares of those brands. In any counterfactual, brand b's owner will be denoted f'. The counterfactual market share of firm f' is

$$S'_{f'n} = \sum_{b \in \mathcal{F}'_{f'}} \mathbb{I}_{bn} \hat{s}_{bn} s_{bn} \tag{22}$$

Since in our model equilibrium markups are directly related to the firm-destination market share, we can use equation (11) to compute the proportional change in the Lerner index for each brand under the two alternative conduct assumptions:

$$\underbrace{\hat{L}_{bn} = \frac{\sigma - (\sigma - \eta) S_{fn}}{\sigma - (\sigma - \eta) S'_{f'n}}}_{\text{Bertrand}} \quad \text{and} \quad \underbrace{\hat{L}_{bn} = \frac{1 + (\sigma/\eta - 1) S'_{f'n}}{1 + (\sigma/\eta - 1) S_{fn}}}_{\text{Cournot}}. \tag{23}$$

Note that brands which remain under the same ownership have $f' = f$ and these equations can still be applied. Since $\mu = 1/(1 - L)$, the brand's price-cost markup adjusts as follows:

$$\hat{\mu}_{bn} = \frac{1 - L_{bn}}{1 - \hat{L}_{bn} L_{bn}}. \tag{24}$$

With these markup adjustments calculated, we can compute the brand-level market share changes. The main cause of brand-level market share changes is the adjustment of markups resulting from the change in ownership. However, the method allows for changes in the cost-adjusted appeal of brand b to market n, denoted $\hat{\varphi}_{bn}$. These could enter through two channels. First, a brand with a new owner f' inherits the potentially different $\varphi^F_{f'}$. Second, if $h(f') \neq h(f)$ then headquarters frictions, δ^F, change.

The proportional change in brand-level market share is given by

$$\hat{s}_{bn} = \left(\frac{\hat{\mu}_{bn}}{\hat{\varphi}_{bn} \hat{P}_{gn}} \right)^{1-\sigma} \quad \text{with} \quad \hat{P}_{gn} = \left(\sum_k \mathbb{I}_{kn} s_{kn} (\hat{\mu}_{kn}/\hat{\varphi}_{kn})^{1-\sigma} \right)^{\frac{1}{1-\sigma}}. \tag{25}$$

The resulting $\hat{s}_{bn}$ is the same as the one obtained by solving for the equilibrium, s_{bn}, before and after the friction change and taking the ratio. The advantage is that it can be calculated without knowing the *levels* of all the model's parameters. Compared to the existing EHA methods covered by Costinot and Rodriguez-Clare (2014), the innovation is to account for endogenous markup adjustment ($\hat{\mu}_{bn}$ in equation 25).

The summation in equation (25) includes fringe brands whose individual market shares are not observed.[37] Since the fringe brands are monopolistically competitive, their markups are fixed at $\sigma/(\sigma-1)$, implying $\hat{\mu}_{0n} = 1$. The counterfactuals hold ownership constant in the fringe and hence also hold their φ constant. Therefore, the aggregate market share of the fringe brands—which we do observe and denote as S_{0n}—is all we need to complete the counterfactual price index adjustment:

$$\hat{P}_{gn} = \left(S_{0n} + \sum_{k \in \text{listed}_n} \mathbb{I}_{kn} s_{kn} (\hat{\mu}_{kn}/\hat{\varphi}_{kn})^{1-\sigma} \right)^{\frac{1}{1-\sigma}}. \tag{26}$$

For the listed brands, markups and brand type adjust in response to changes in ownership, leading to new market shares determined by equation (25). The share of fringe brands evolves according to $\hat{S}_{0n} = \hat{P}_{gn}^{\sigma-1}$.

Finally, we need to account for the consequences of the counterfactual shock at the upper level. The assumption that each sector is too small to affect the aggregate price index implies that $\hat{P}_n = 1$ and $\hat{X}_n = 1$. Hence, expenditures in category g adjust to price

[37] Redding and Weinstein (2018) address an analogous problem, showing how to construct a CES price index with only aggregate information on the share of expenditure on non-traded varieties.

changes according to $\hat{X}_{gn} = \hat{P}_{gn}^{1-\eta}$.

The algorithm for computing the counterfactual scenario involves the following steps:

0. Initialize s_{bn} and S_{fn} at their historical levels and set $\hat{s}_{bn} = 1, \forall bn$.
1. Apply equation (22) with the new ownership sets, $\mathcal{F}'_{f'}$, to obtain the vector of counterfactual firm-destination market shares, $S'_{f'n}$.
2. Calculate the conduct-specific vector of changes in markups, $\hat{\mu}_{bn}$, applying (23) and (24) to the current values of S_{fn} and $S'_{f'n}$.
3. Inputting $\hat{\varphi}_{bn}$ and the new $\hat{\mu}_{bn}$ into (25), handling the fringe as specified in (26), calculate the new brand market share adjustments $\hat{s}_{bn}$. Go to step 1.

Repeat steps 1–3 until the vector of $\hat{s}_{bn}$ stops changing.

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 (26). 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.[38]

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.[39]

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

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

[39] The ABI/Modelo decision by US DOJ and European Commission decision (Case M.7881: AB INBEV/SABMILLER) on the SABMiller acquisition point to both effects to justify divestitures.

Table 8: What if antitrust authorities had been more permissive?

Country	%Chg. P_{gnt} ($\hat{\varphi}_{bn} = 1$) Bertrand	%Chg. P_{gnt} ($\hat{\varphi}_{bn} = 1$) Cournot	%Chg. P_{gnt} ($\hat{\varphi}_{bn} \neq 1$) Bertrand	%Chg. P_{gnt} ($\hat{\varphi}_{bn} \neq 1$) Cournot
United States	4.23	5.88	4.37	5.90
United Arab Emirates	1.13	1.91	1.12	1.87
Netherlands	1.04	2.04	0.08	0.99
Hungary	1.03	1.83	-0.37	0.11
Italy	0.79	1.58	0.05	0.74
Czechia	0.54	0.78	-1.76	-1.91
Slovakia	0.20	0.34	-1.58	-1.81
Poland	0.00	0.00	-1.72	-2.07

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%.

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 one percent are displayed in Table 8. 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 φ_{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 C.8).

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 four to six percent higher. The highest price increase occurs under Cournot competition. The third and fourth columns show that taking into account changes in φ_{bn} leads to a small exacerbation 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.[40]

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 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 leading 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 or fully 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.[41] However, the move of HQ from Belgium to Japan increases frictions by enough to raise the price index of beer by 2.2 to 2.3%. 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 one to six percent in five 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.

[40]Non-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.

[41]In 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.

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 9 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 market presence in the markets where these brands had high market shares, this 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.

Table 9: What if antitrust authorities had followed EU/US lead?

Country	%Chg. P_{gnt} ($\hat{\varphi}_{bn} = 1$)		%Chg. P_{gnt} ($\hat{\varphi}_{bn} \neq 1$)	
	Bertrand	Cournot	Bertrand	Cournot
Colombia	-30.21	-25.87	-29.59	-24.62
Ecuador	-25.26	-22.69	-24.68	-21.49
Peru	-19.59	-14.05	-19.14	-12.62
Uruguay	-10.12	-11.54	-10.51	-11.73
Dominican Republic	-7.05	-4.18	-7.34	-4.27
Canada	-2.65	-5.50	-2.09	-4.58
Argentina	-2.24	-4.22	-2.25	-4.11
Australia	-1.97	-4.32	-3.92	-5.94
United Arab Emirates	-1.72	-3.77	-1.65	-3.47
Bolivia	-1.63	-2.12	-1.72	-2.17
Mexico	-1.35	-2.94	-2.06	-3.27
Chile	-1.16	-2.71	-1.33	-2.76
South Africa	-1.11	-2.05	-2.86	-3.48
Guatemala	-0.66	-1.50	-0.79	-1.58
India	-0.37	-0.95	-1.56	-2.01

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 SABMiller had acquired the local star brands. Forcing divestitures would have reduced the beer price index by 14–30% depending on the country and assumptions. The Dominican Republic and Uruguay would also experience gains as large, or larger, than those generated by divestiture for the US. For all countries except the first three listed in table 9, forcing divestiture yields larger price reductions under Cournot conduct than Bertrand. The intuition for why the Bertrand effects are stronger for Colombia, Ecuador, Peru can be found in the convexity of the Lerner index as a function of market share under Bertrand conduct (shown in Figure 3(a)). Those three countries started out in the region of market shares where further consolidation boosts markups more under Bertrand.

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 9 suggests that implementing the three divestitures (Labatt, Modelo, and SABMiller EU brands) would have saved Canadian consumers between 2.7% and 6.4%. Australian beer drinkers would gain 1.9% to 4.3%. Mexico could also have generated substantial gains through compelling divestiture of the Modelo brands in the Mexican market.

The price reductions reported in Table 9 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.[42] 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.

Table 10 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 points or more "will be presumed likely to enhance market power."[43] Table 10 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 0.2–4.1% higher for the average country.[44] The biggest increase is

[42]The 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).

[43]`https://www.justice.gov/atr/horizontal-merger-guidelines-08192010`

[44]Most of the average price increases are smaller than 4% average that Kwoka (2014) obtained in a meta-analysis of 47 merger retrospectives covering a variety of different products.

Table 10: Summary of outcomes of the counterfactual restoring 2007 brand owners

Category	# of Countries	Conduct assumed	Chg. HHI Mean	Chg. HHI Median	%Chg. P_{gnt} Mean	%Chg. P_{gnt} Median
				with $\hat{\varphi}_{bn} = 1$		
Beer	76	Bertrand	424	212	2.41	0.68
Beer	76	Cournot	481	251	3.10	1.56
Spirits	75	Bertrand	67	20	0.22	0.05
Spirits	75	Cournot	67	21	0.38	0.10
				with $\hat{\varphi}_{bn} \neq 1$		
Beer	76	Bertrand	376	172	3.30	1.16
Beer	76	Cournot	428	215	4.09	1.86
Spirits	75	Bertrand	48	18	0.87	0.41
Spirits	75	Cournot	47	20	1.02	0.51

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.

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).

Appendix D graphs the counterfactual concentration and price index changes for all countries in our data set. The counterfactual points to sizeable price increases in just two spirits markets: Turkey and Tunisia. In the former, Diageo's acquisition of the owner of Yeni Raki, the most popular spirit in the country, leads to a price rise between 3% (Bertrand, $\hat{\varphi}_{bn} = 1$) and 10% (Cournot, 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). Since Absolut's prior owner was also foreign and had a similar group fixed effect, the acquisition did not change $\hat{\varphi}_{bn}$ by much. The market power effect raises the Tunisian spirit price index by 3–4%.

Figures D.1(a) and D.1(c) illustrate how pure market power effects vary with concentration (holding brand type constant, $\hat{\varphi}_{bn} = 1$). For all countries where the rise in concentration in less than 1000, the rise in the price index is roughly linear in the change in the HHI. This corroborates the local approximation result in Proposition 5 of Nocke and Schutz (2018a). For beer, we see some non-linearity for HHI changes over 1000 under Bertrand, but linearity is a good approximation globally for Cournot. Figures D.1(b) and D.1(d) build in changes in φ_{bn} resulting from owner and HQ changes. The positive relationship between counterfactual changes in price index and concentration persists, but departs considerably from the tight line for spirits.

The model can be used to calculate changes in markups over markets to construct the counterfactual change in each firm's global Lerner index: ΔL_f. This provides a perspective on how mergers have transformed firms' market power which is complementary to the market-level perspective captured by changes in concentration and price indexes. The consolidated markup L_f, depends on brand-level market shares and the way they map to owners. As can be seen from inspecting equation (12), L_f is high when the firm has high market share in the markets that contribute importantly to its global revenues.

Figure 6: Effects of ownership changes 2007–18 on firm-level markups

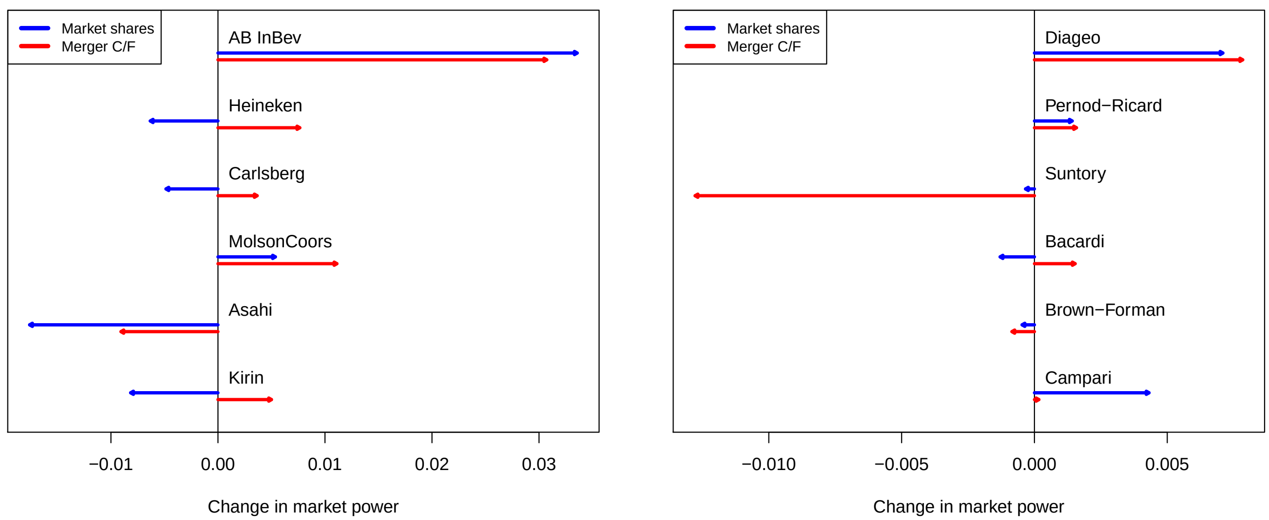

(a) Beer (b) Spirits

We start by calculating the change in consolidated markups implied by the historical evolution of market shares. That is, we calculate ΔL_f by plugging the factual market shares and ownership patterns into equation (12) using 2007 and 2018 data. Let $\mathbf{S}_{ft}$ denote the vector of all firm-level market shares in each country n in year t. Recalling that $o(b,t)$ gives the mapping of brands to owners in any year t, the arrows shown in blue in Figure 6 correspond to

$$\Delta L_f^{\text{blue}} = L_f(\mathbf{S}_{f18}, o(b,18)) - L_f(\mathbf{S}_{f07}, o(b,07)).$$

Next, we calculate the change in consolidated markups that isolates those changes coming purely from ownership changes entailed in switching from $o(b,07)$ to $o(b,18)$:

$$\Delta L_f^{\text{red}} = L_f(\mathbf{S}_{f18}, o(b,18)) - L_f(\mathbf{S}'_{f18}, o(b,07)),$$

where $\mathbf{S}'_{f18}$ is the counterfactual vector of market shares if the brand owners of 2007 repossessed their holdings in that year. Figure 6 shows this with red arrows.

The first terms in both the blue and red versions are the same, but the subtracted terms differ. The blue ΔL_f subtracts historical market shares from 2007, whereas the red ΔL_f subtracts a counterfactual based on 2018 market shares and 2007 ownership. The ΔL_f^{blue} arrows in Figure 6 combine changes in firm markups coming from altering brand portfolios with changes in brand type of each firm's *incumbent* brands in each market. The ΔL_f^{red} arrows exclude the φ_{bn} changes in incumbent brands.

The most important takeaway from Figure 6 is the very close match between ΔL_f^{blue} and ΔL_f^{red} for the largest firms in each category, AB InBev and Diageo, which account for 26% and 10% of the world beer and spirits markets. M&A essentially tells the whole story for the growth in markups for these multinationals. Brand performance was static but, by combining brands to increase firm-level market share, these two firms increased their aggregate worldwide market power. AB InBev benefits a little from changes in market share by incumbent brands—a kind of superstar effect at the brand level. Diageo, on the other hand, is held back by subpar incumbent brand performance.

The second and third largest beers makers, Heineken and Carlsberg, present a puzzle in that their M&A activities should have been *increasing* market power, but the actual evolution of historical market shares points to falling market power. The explanation is that, despite numerous acquisitions in multiple markets (as shown in Figure 1), the losses of market share for flagship brands in the strongholds of those two firms (notably Spain, Poland, and Greece for Heineken, and all Nordic countries for Carlsberg) dominated the gains in markets entered via acquisitions. This resulted in $\Delta L_f^{\text{blue}} < 0$ and $\Delta L_f^{\text{red}} > 0$ for both firms.

Asahi and Kirin represent paradoxical cases of firms whose expansion abroad led to *lower* indexes of market power. This happens because their portfolios transformed from a complete Japan focus, where their market shares were dominant (40% and 31%, respectively), to diversified positions where lower market share brands contribute substantially to total sales.[45] In the case of Asahi, this is the primary reason for its decline in L_f over the decade. Kirin, however, suffered from the same incumbent brand decline experienced by Heineken and Carlsberg.

In the case of spirits, we see one case, Suntory, where M&A dragged down the firm-level measure of market power. This was not because Suntory was selling off brands,

[45] Both firms obtain 98% of sales from Japan at the start of our sample but, by 2018, Japan's weight falls to 54% and 62%. In the new markets, the firms acquired strong brands but they only rarely matched their Japan market shares.

but rather because the brands it gained gave it higher sales shares in markets where it had low L_{fn}. Before purchasing Beam, Suntory had high market share (16%) in Japan and negligible sales elsewhere. With Beam's brands, Suntory's sales in the US vaulted over their sales at home. However, the Beam brands captured only an 8% share of the US market, implying relatively low markups. This depressed Suntory's worldwide L_f by over a percentage point.

The one firm in Figure 12 that displays superstar effects is Campari. This is in large part attributable to the outstanding growth of one of its incumbent brands, Aperol. The φ_{bn} of this brand rises in several major markets. The parent company also started to offer the brand in 21 new markets.

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 just 2–7% 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 a 10–20% penalty on cost-adjusted appeal from foreign acquisitions with little in the way of predictable efficiencies. From the firm's point of view, there may be compensating reductions in fixed costs, but the methods we use here cannot recover such effects. The other potential benefit to firms is increased market power, a concern our counterfactuals show to be important—but highly heterogeneous 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.

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 consumer prices increases of 20–30% could have been avoided.

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.

Arkolakis, C., A. Costinot, D. Donaldson, and A. Rodríguez-Clare (2018). The Elusive Pro-Competitive Effects of Trade. *The Review of Economic Studies 86*(1), 46–80.

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. 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.

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., 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.

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.

Feenstra, R. C. (2003). *Advanced International Trade: Theory and Evidence*. Princeton University Press.

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.

Grassi, B. (2017). Io in i-o: Size, industrial organization, and the input-output network make a firm structurally important. mimeo.

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 (2019a). Brands in motion: How frictions shape multinational production. *American Economic Review 109*(9), 3073–3124.

Head, K. and T. Mayer (2019b). Poor substitutes? Counterfactual methods in IO and trade compared. 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., M. J. Melitz, and G. I. Ottaviano (2014). Market size, competition, and the product mix of exporters. *American Economic Review 104*(2), 495–536.

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., G. Sheu, and M. Weinberg (2019). Oligopolistic price leadership and mergers: The united states beer industry.

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

Miravete, E. J., K. Seim, and J. Thurk (2018). Market power and the Laffer curve? mimeo.

Neary, J. P. (2016). International trade in general oligopolistic equilibrium. *Review of International Economics 24*(4), 669–698.

Nocke, V. and N. Schutz (2018a). An aggregative games approach to merger analysis in multiproduct-firm oligopoly. Discussion Paper 12905, CEPR.

Nocke, V. and N. Schutz (2018b). 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.

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

Syverson, C. (2019b). Macroeconomics and market power: Facts, potential explanations and open questions. *Brookings Economic Studies Report*.

World Bank, T. (2020). *World Development Report 2020: Trading for development in the age of global value chains*. World Bank Publications.

Appendix

A 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 A.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.[46]

Table A.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

[46] Bernard 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 A.1: Global giants are rare

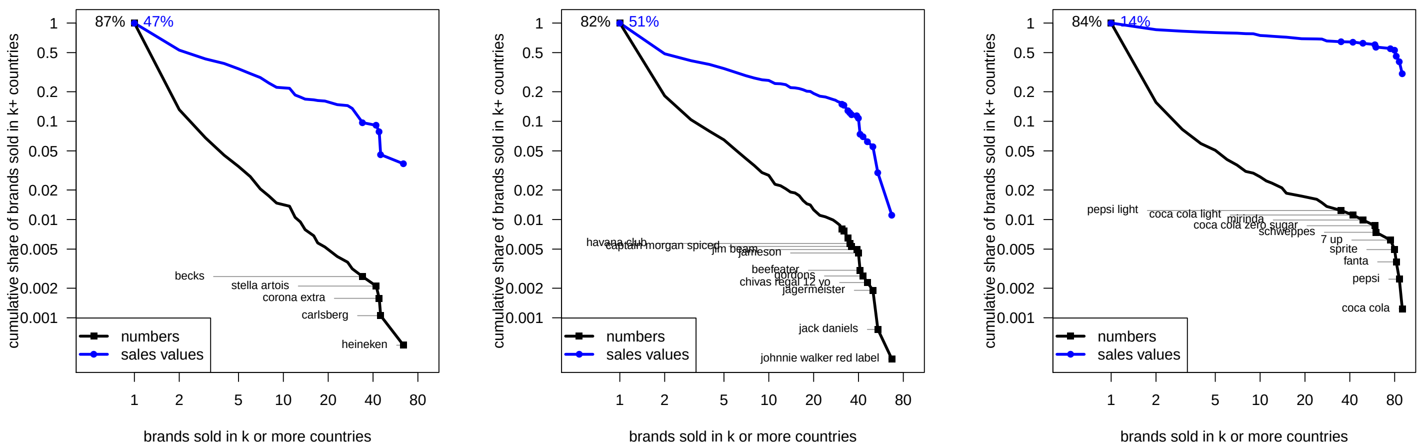

(a) Beer (b) Spirits (c) Carbonates

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.

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

Sample frame	Add rate (in percent)	Drop rate (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 market:		
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:		
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 market:		
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

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 A.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 columnn 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 A.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 A.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

Figure A.2: Small changes in brand offerings following ownership changes

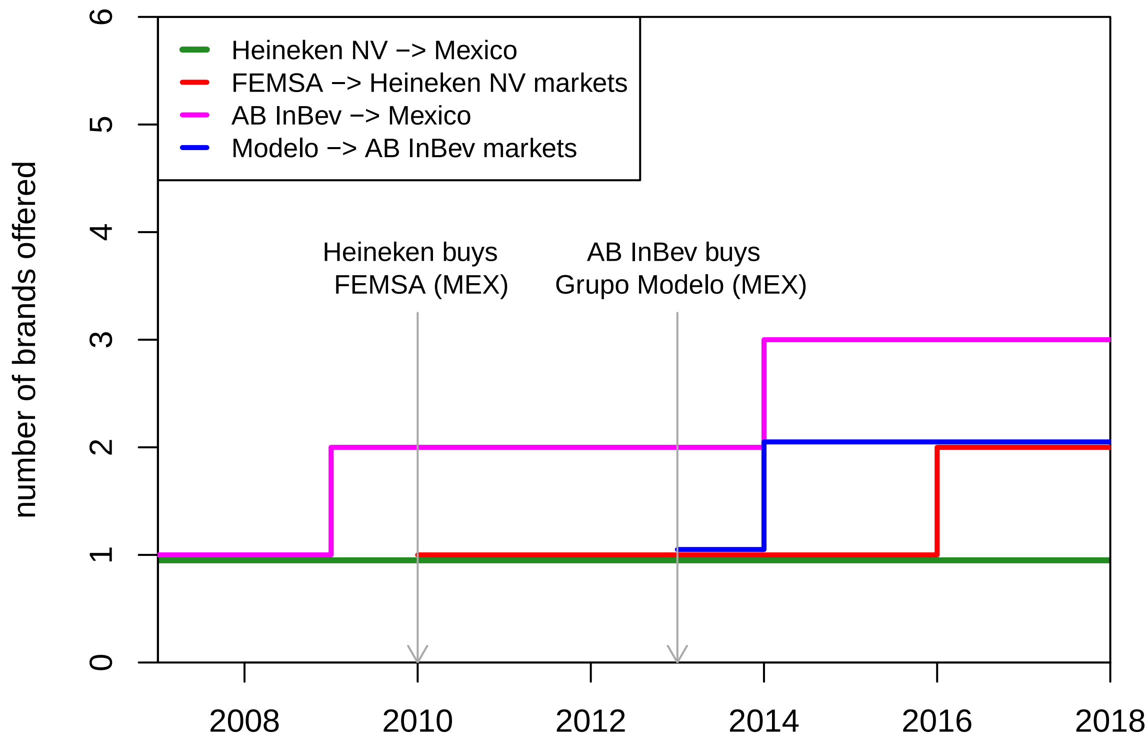

(a) Acquisitions of the two largest Mexican breweries

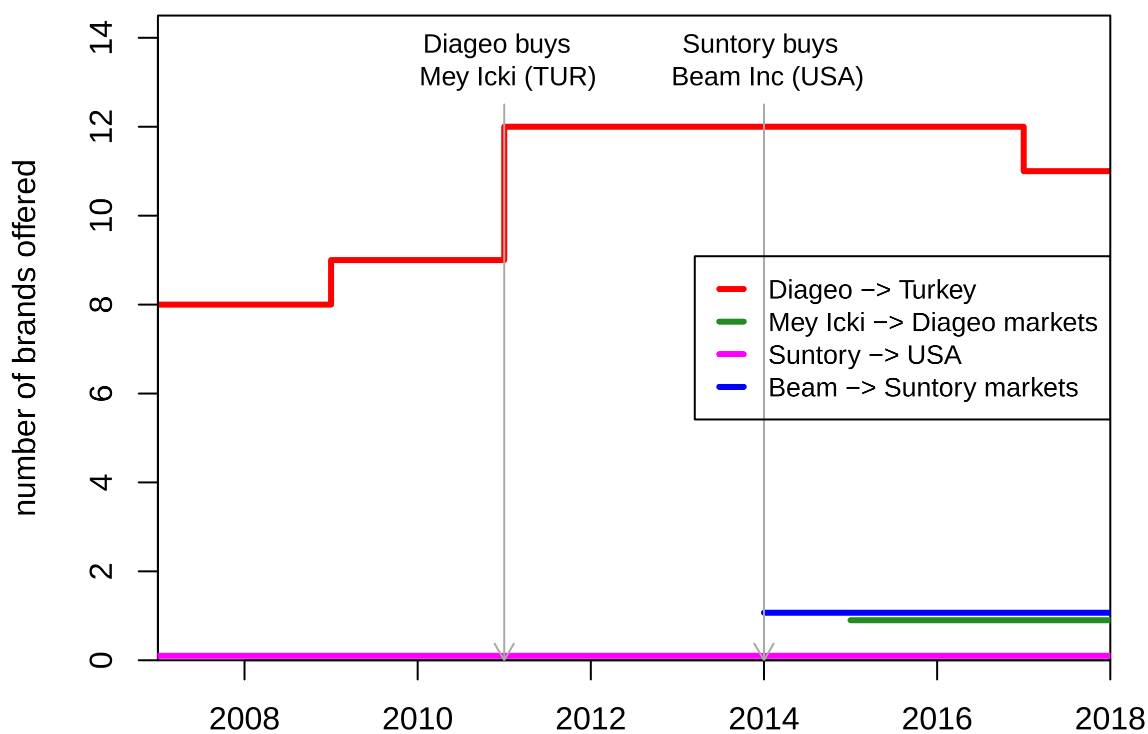

(b) Diageo and Suntory purchases of Mey Icki and Beam Inc.

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 A.1.

B Connectivity of the brand-firm network

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

Product group	# Firms		Mobility		Sales share	
Beer	90	21	13.4	50.1	80.0	70.8
Spirits	93	17	8.0	32.5	57.5	41.9
Wine	12	2	6.4	27.5	6.3	2.9
Water	68	3	2.3	11.3	58.9	43.4
Carbonates	44	4	3.3	11.5	91.2	65.7
Juice	60	2	2.7	13.0	44.5	2.8
Coffee	3	NA	2.7	NA	33.1	NA
≥ 10 movers		✓		✓		✓

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.

In the third and fourth columns of Table B.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 B.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.

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

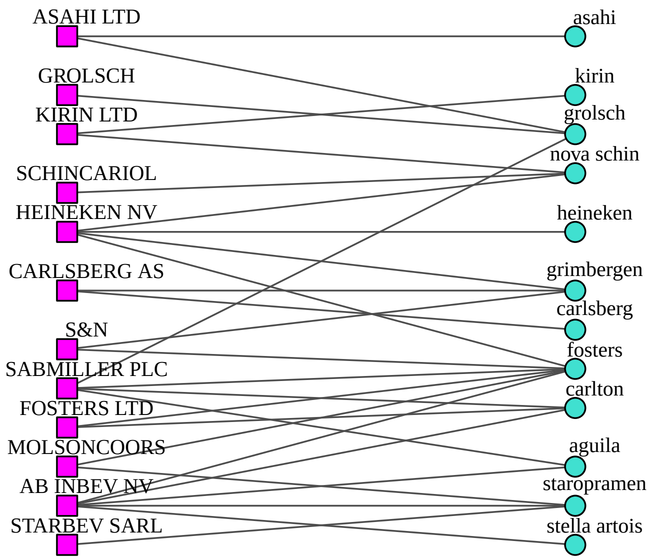

(a) A connected set of firms and brands

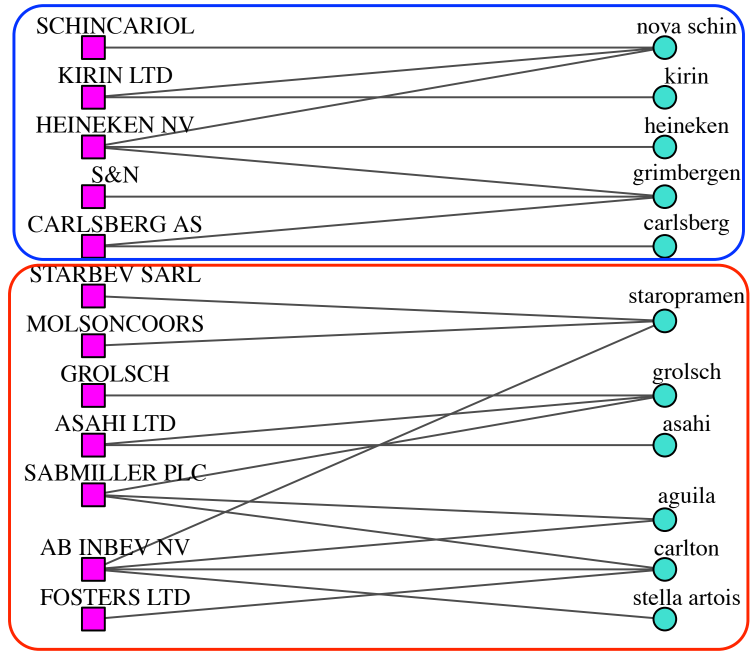

(b) Without the Fosters brand, the sets disconnect

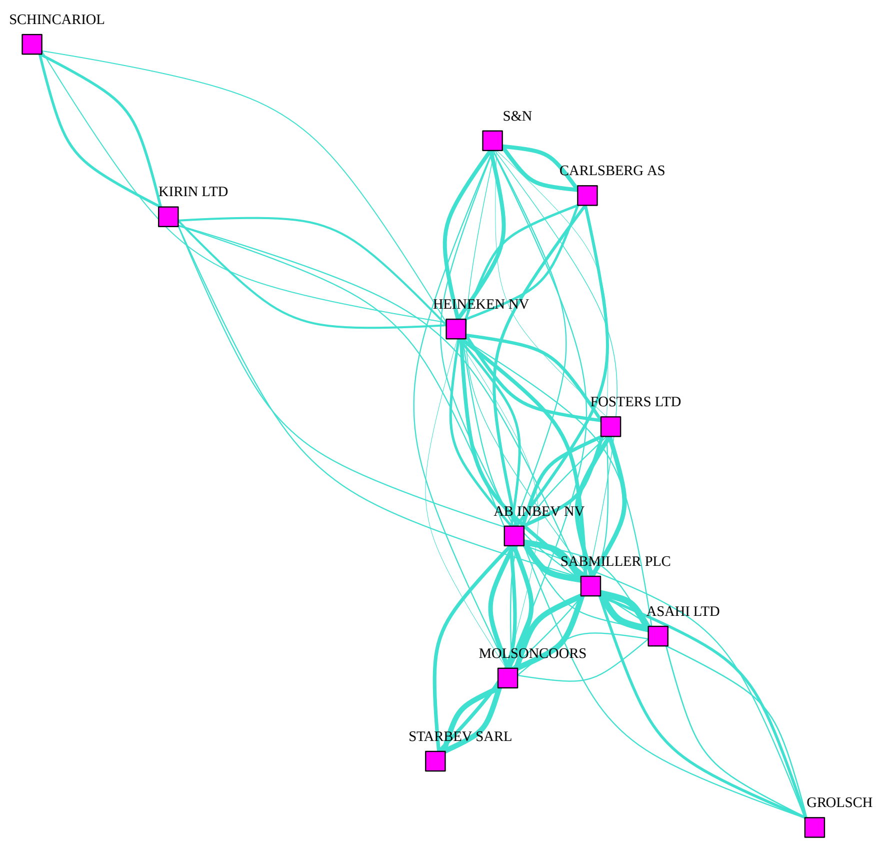

(c) The *induced* firm-to-firm network from panel (a)

(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 λ_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 λ_2. Chung (1997) shows that the maximum λ_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, $\lambda_2 \to 1$.

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

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

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

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

Figure B.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).

C Additional regression results

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

			Bertrand		Cournot	
	$\ln s_{bn}$	$\ln A_{bn}$	$\ln \varphi_{bn}$	$\ln \mu_{bn}$	$\ln \varphi_{bn}$	$\ln \mu_{bn}$
home	1.020^{a}	0.211^{a}	0.355^{a}	0.022^{a}	0.375^{a}	0.041^{a}
	(0.127)	(0.068)	(0.048)	(0.004)	(0.050)	(0.007)
distance	-0.110^{a}	0.030	-0.040^{a}	-0.001	-0.041^{a}	-0.003^{c}
	(0.035)	(0.020)	(0.014)	(0.001)	(0.014)	(0.002)
common language	0.053	-0.053	0.011	0.0004	0.012	0.001
	(0.076)	(0.049)	(0.030)	(0.002)	(0.031)	(0.003)
home (HQ)	0.285^{a}	0.082	0.140^{a}	0.018^{a}	0.154^{a}	0.032^{a}
	(0.090)	(0.051)	(0.036)	(0.003)	(0.037)	(0.006)
distance (HQ)	0.006	0.009	0.007	-0.0005	0.006	-0.002
	(0.026)	(0.016)	(0.011)	(0.001)	(0.011)	(0.001)
com. lang. (HQ)	0.096^{c}	0.046	0.042^{c}	0.001	0.044^{c}	0.003
	(0.058)	(0.035)	(0.023)	(0.003)	(0.024)	(0.004)
Observations	95,299	95,299	95,299	95,299	95,299	95,299
R^2	0.651	0.649	0.589	0.891	0.596	0.846

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).

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

			Bertrand		Cournot	
	$\ln s_{bn}$	$\ln A_{bn}$	$\ln \varphi_{bn}$	$\ln \mu_{bn}$	$\ln \varphi_{bn}$	$\ln \mu_{bn}$
home	1.056^a	0.238^a	0.363^a	0.019^a	0.379^a	0.035^a
	(0.151)	(0.079)	(0.056)	(0.004)	(0.057)	(0.007)
distance	-0.083^b	0.049^b	-0.026^c	-0.001	-0.027^c	-0.002
	(0.038)	(0.023)	(0.015)	(0.001)	(0.015)	(0.002)
common language	0.051	-0.051	0.012	0.001	0.012	0.001
	(0.079)	(0.052)	(0.032)	(0.002)	(0.032)	(0.003)
home (HQ)	0.263^b	0.084	0.154^a	0.041^a	0.186^a	0.073^a
	(0.117)	(0.065)	(0.046)	(0.005)	(0.048)	(0.008)
distance (HQ)	0.033	0.010	0.017	0.001	0.016	0.0001
	(0.035)	(0.021)	(0.014)	(0.001)	(0.014)	(0.002)
com. lang. (HQ)	0.117^c	0.054	0.054^b	0.004	0.058^b	0.008
	(0.066)	(0.041)	(0.027)	(0.003)	(0.028)	(0.005)
Observations	64,968	64,968	64,968	64,968	64,968	64,968
R^2	0.598	0.568	0.519	0.876	0.527	0.827

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).

C.1 Results pooling 7 Beverages

Table C.3: Pooled regressions, 7 beverages, with firm fixed effects

			Bertrand		Cournot	
	$\ln s_{bn}$	$\ln A_{bn}$	$\ln \varphi_{bn}$	$\ln \mu_{bn}$	$\ln \varphi_{bn}$	$\ln \mu_{bn}$
home	1.004^a	0.269^a	0.395^a	0.018^a	0.409^a	0.032^a
	(0.099)	(0.055)	(0.040)	(0.003)	(0.041)	(0.005)
distance	-0.155^a	0.002	-0.052^a	-0.0003	-0.053^a	-0.001
	(0.029)	(0.017)	(0.012)	(0.001)	(0.012)	(0.001)
common language	0.117^c	-0.022	0.038	0.001	0.039	0.003
	(0.063)	(0.038)	(0.026)	(0.002)	(0.026)	(0.003)
home (HQ)	0.381^a	0.147^a	0.177^a	0.020^a	0.194^a	0.037^a
	(0.080)	(0.044)	(0.033)	(0.003)	(0.034)	(0.005)
distance (HQ)	0.026	0.017	0.013	-0.001	0.011	-0.002^c
	(0.026)	(0.015)	(0.011)	(0.001)	(0.011)	(0.001)
com. lang. (HQ)	0.152^a	0.062^b	0.068^a	0.003	0.070^a	0.006
	(0.053)	(0.032)	(0.022)	(0.002)	(0.023)	(0.004)
Observations	170,578	170,578	170,578	170,578	170,578	170,578
R^2	0.735	0.699	0.667	0.941	0.672	0.912

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).

Table C.4: Pooled regressions, 7 beverages, without firm fixed effects

			Bertrand		Cournot	
	$\ln s_{bn}$	$\ln A_{bn}$	$\ln \varphi_{bn}$	$\ln \mu_{bn}$	$\ln \varphi_{bn}$	$\ln \mu_{bn}$
home	1.023^{a}	0.282^{a}	0.409^{a}	0.023^{a}	0.426^{a}	0.040^{a}
	(0.093)	(0.052)	(0.038)	(0.003)	(0.039)	(0.005)
distance	-0.148^{a}	-0.005	-0.052^{a}	0.00004	-0.053^{a}	-0.001
	(0.027)	(0.015)	(0.011)	(0.001)	(0.011)	(0.001)
common language	0.125^{b}	-0.019	0.041^{c}	0.0002	0.042^{c}	0.001
	(0.061)	(0.036)	(0.025)	(0.002)	(0.025)	(0.003)
home (HQ)	0.286^{a}	0.093^{b}	0.125^{a}	0.010^{a}	0.134^{a}	0.020^{a}
	(0.069)	(0.037)	(0.029)	(0.002)	(0.030)	(0.004)
distance (HQ)	0.011	0.016	0.008	-0.001	0.007	-0.002
	(0.021)	(0.012)	(0.009)	(0.001)	(0.009)	(0.001)
com. lang. (HQ)	0.109^{b}	0.043	0.048^{b}	0.003	0.050^{b}	0.005
	(0.050)	(0.029)	(0.021)	(0.002)	(0.021)	(0.003)
Observations	170,578	170,578	170,578	170,578	170,578	170,578
R^2	0.726	0.689	0.653	0.935	0.658	0.901

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).

C.2 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 C.5 shows fixed effect correlations for regressions on all firms in the largest connected set. The underlying regressions in table C.6 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 C.6 show the sign and magnitude of assortative matching.

Table C.5: Correlations between fixed effects in the largest connected set

	Brand				Firm			
Dep. var.:	share (s_{bn})	appeal (A_{bn})	type B (φ_{bn})	type C (φ_{bn})	share (s_{bn})	appeal (A_{bn})	type B (φ_{bn})	type C (φ_{bn})
Beer								
brand market share	1.278							
brand appeal	0.749	1.191						
brand type B	0.991	0.738	1.313					
brand type C	0.985	0.736	0.998	1.243				
firm market share	-0.538	-0.250	-0.510	-0.508	0.385			
firm appeal	-0.352	-0.367	-0.310	-0.312	0.688	0.266		
firm type B	-0.521	-0.229	-0.497	-0.496	0.992	0.662	0.359	
firm type C	-0.499	-0.214	-0.475	-0.477	0.980	0.659	0.995	0.340
Spirits								
brand market share	0.621							
brand appeal	0.719	0.630						
brand type B	0.999	0.718	0.644					
brand type C	0.998	0.717	1.000	0.636				
firm market share	-0.511	-0.266	-0.514	-0.515	0.216			
firm appeal	-0.376	-0.400	-0.382	-0.386	0.681	0.102		
firm type B	-0.496	-0.257	-0.500	-0.501	0.997	0.688	0.231	
firm type C	-0.482	-0.249	-0.486	-0.488	0.990	0.691	0.998	0.236

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 C.6, the correlations become much smaller, and not even systemat-

Table C.6: Correlations between fixed effects in the AGSU restricted sample

	Brand				Firm			
Dep. var.:	share (s_{bn})	appeal (A_{bn})	type B (φ_{bn})	type C (φ_{bn})	share (s_{bn})	appeal (A_{bn})	type B (φ_{bn})	type C (φ_{bn})
Beer								
brand market share	0.937							
brand appeal	0.780	1.124						
brand type B	0.989	0.768	0.988					
brand type C	0.982	0.768	0.998	0.951				
firm market share	-0.100	-0.141	-0.099	-0.108	0.036			
firm appeal	-0.055	-0.158	-0.056	-0.070	0.886	0.070		
firm type B	-0.071	-0.101	-0.069	-0.076	0.967	0.846	0.036	
firm type C	-0.036	-0.053	-0.035	-0.040	0.910	0.780	0.981	0.036
Spirits								
brand market share	0.394							
brand appeal	0.719	0.407						
brand type B	0.999	0.714	0.400					
brand type C	0.997	0.711	1.000	0.393				
firm market share	-0.030	0.088	-0.039	-0.044	0.057			
firm appeal	-0.108	-0.021	-0.126	-0.139	0.720	0.044		
firm type B	-0.027	0.092	-0.035	-0.041	0.991	0.714	0.065	
firm type C	-0.021	0.094	-0.029	-0.035	0.976	0.704	0.996	0.070

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 C.6.

Table C.7: The explanatory power of owner fixed effects: Cournot conduct

Type of FE	# of FE	λ_2	ΔR^2	Varshr	FE Corr
		Beer			
Firms (All)	464	0.000	0.007	NA	NA
Firms (Largest connected set, AKM)	90	0.013	0.009	0.340	-0.477
Firms (Leave-one-out, KSS)	49	0.071	0.005	0.057	-0.080
Firms (High mobility, AGSU)	22	0.171	0.005	0.036	-0.040
Clusters (BLM)	15	0.461	0.001	0.034	0.159
Clusters (BLM)	10	0.548	0.001	0.038	0.212
Clusters (BLM)	5	0.618	0.001	0.033	0.222
		Spirits			
Firms (All)	849	0.000	0.007	NA	NA
Firms (Largest connected set, AKM)	93	0.013	0.007	0.236	-0.488
Firms (Leave-one-out, KSS)	41	0.010	0.015	0.108	-0.144
Firms (High mobility, AGSU)	18	0.071	0.007	0.070	-0.035
Clusters (BLM)	15	0.426	0.002	0.054	0.169
Clusters (BLM)	10	0.436	0.002	0.058	0.175
Clusters (BLM)	5	0.904	0.001	0.024	0.299

Notes: # of FE is either number of firms or clusters. λ_2 measures network connectivity. ΔR^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.

Table C.8: Friction estimates, alternative heterogeneity assumptions: Cournot conduct

	Beer			Spirits		
Fixed effects:	$b+f$	$b+k$	bf	$b+f$	$b+k$	bf
home	0.469^{a}	0.497^{a}	0.477^{a}	0.281^{a}	0.273^{a}	0.279^{a}
	(0.057)	(0.056)	(0.058)	(0.068)	(0.066)	(0.069)
distance	-0.077^{a}	-0.068^{a}	-0.086^{a}	-0.032^{c}	-0.031^{c}	-0.032^{c}
	(0.019)	(0.018)	(0.020)	(0.019)	(0.019)	(0.019)
common language	0.092^{b}	0.109^{a}	0.086^{b}	-0.019	-0.018	-0.020
	(0.042)	(0.041)	(0.042)	(0.040)	(0.039)	(0.040)
home (HQ)	0.136^{b}	0.082^{c}	0.128^{b}	0.232^{a}	0.219^{a}	0.248^{a}
	(0.056)	(0.047)	(0.059)	(0.058)	(0.053)	(0.061)
distance (HQ)	-0.038^{b}	-0.037^{a}	-0.037^{c}	0.030^{c}	0.029^{c}	0.031^{c}
	(0.017)	(0.013)	(0.020)	(0.017)	(0.015)	(0.018)
com. lang. (HQ)	-0.022	-0.038	-0.011	0.079^{b}	0.072^{b}	0.079^{b}
	(0.038)	(0.035)	(0.041)	(0.031)	(0.030)	(0.032)
Observations	34,675	34,675	34,675	60,624	60,624	60,624
R^2	0.744	0.737	0.756	0.553	0.547	0.557
RMSE	0.245	0.246	0.241	0.388	0.387	0.385

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).

D Restoring 2007 owners: concentration and price indexes

Figure D.1 illustrates the model-based quantification of the impact of mergers and acquisitions occurring over the decade after 2007. The graphs in the left column hold φ_{bn} constant wheres the graphs on the right use our HQ friction estimates and group fixed effect to capture changes in φ_{bn}. The vertical axes display changes in the price index attributed to the 2007–2018 ownership changes. The horizontal axes show changes in the Herfindahl concentration: $\Delta H = H_{2018} - H'_{2007}$. The upper two panels show results for beer and the lower two panels show the spirits results.

Figure D.1: Counterfactual results: restoring the 2007 owner in 2018

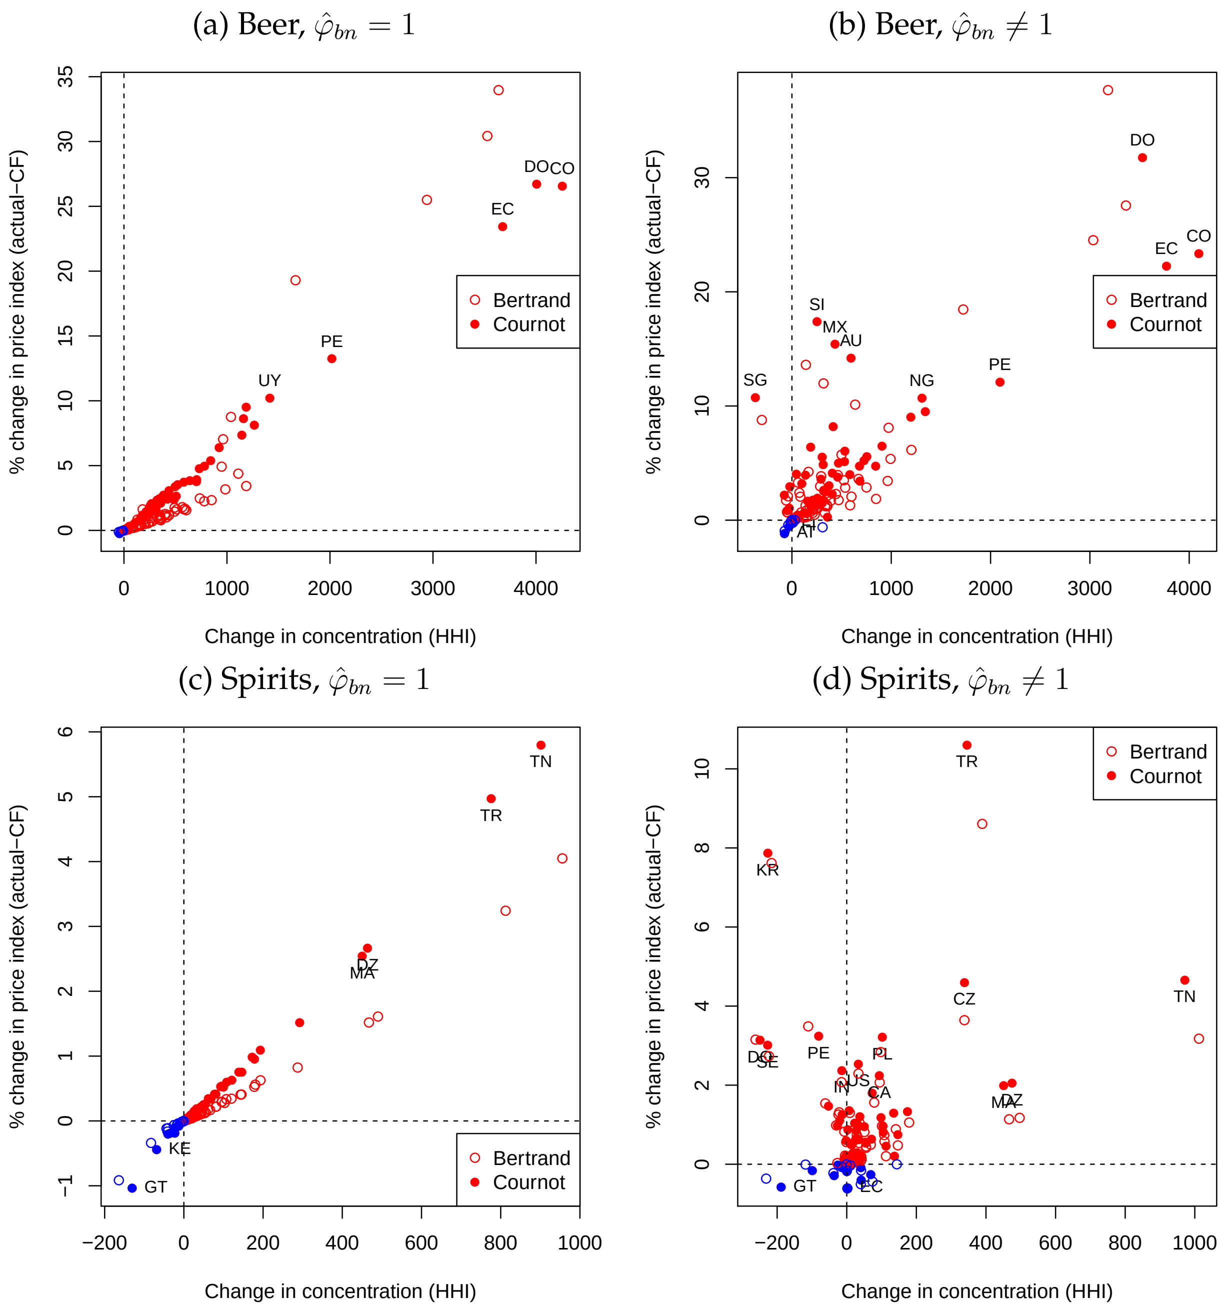

E Concentration and markups

A classical question in industrial organization is how equilibrium markups and overall welfare vary with respect to market concentration, usually measured as a Herfindahl index, that is the sum of squared market shares. In dataset such as ours, we know the aggregate share of the small firms, but not their individual shares. Fringe firms are monopolistically competitive with a Lerner index of $L_0 = 1/\sigma$. The zero mass assumption implies that the Herfindhal index in market n is $H_n = \sum_{f\neq 0} S_{fn}^2$.

The literature specifies and aggregates the markup in several different ways. De Loecker et al. (2020) use a market-share-weighted price to cost ratio. Syverson (2019b) also uses weighted arithmetic means but applies it to the Lerner index. Meanwhile Edmond et al. (2015) and Grassi (2017) use the weighted harmonic mean of μ. We find that for Bertrand competition, the weighted harmonic mean Lerner index gives a neat result whereas for Cournot conduct we can obtain useful results for both the arithmetic and harmonic mean μ. The harmonic mean is signaled with a h superscript, the arithmetic mean with a. For Bertrand competition, recalling that S_{on} is the aggregate market share of "other" firms, we have

$$L_n^h \equiv \left(\sigma S_{on} + \sum_{f\neq o} \frac{S_{fn}}{L_{fn}}\right)^{-1} = \frac{1}{\sigma - (\sigma - \eta) H_n}. \tag{27}$$

As $H_n \to 0$ the aggregate markup goes to the monopolistic competition limit of $L_n^h = 1/\sigma$, whereas sector monopolization ($H_n \to 1$) takes the markup to $L_n^h = 1/\eta$ (the same limiting values we obtain for individual firm Lerner indexes).

Under Cournot the arithmetic mean Lerner index is linear in the Herfindahl,

$$L_n^a \equiv \frac{1}{\sigma} S_{on} + \sum_{f\neq o} S_{fn} L_{fn} = \frac{1}{\sigma} + \left(\frac{1}{\eta} - \frac{1}{\sigma}\right) H_n \tag{28}$$

A special case of this result appears in Syverson (2019b) where he assumes homogeneous goods producers (equivalent to $\sigma \to \infty$) and obtains $L_n^a = H_n/\eta$. Applying the Edmond et al. (2015) definition in the Cournot CES case, the harmonic mean markup is

$$\mu_n^h \equiv \left(\frac{\sigma - 1}{\sigma} S_{on} + \sum_{f\neq o} \frac{S_{fn}}{\mu_{fn}}\right)^{-1} = \left[\frac{\sigma - 1}{\sigma} - \left(\frac{1}{\eta} - \frac{1}{\sigma}\right) H_n\right]^{-1} \tag{29}$$

Now the limiting price-cost ratios are $\mu_n^h = \sigma/(\sigma - 1)$ as $H_n \to 0$ and $\mu_n^h = \eta/(\eta - 1)$ as $H_n \to 1$.[47] The general point is that under both types of conduct, aggregate markups

[47] Burstein et al. (2019) independently derived this relationship and use the fact that $1/\mu^h$ is linear in the

are increasing with the Herfindahl, moving from monopolistic competition to monopoly levels.

De Loecker et al. (2020) use a market share weighted price to cost ratio, that is

$$\mu_n^a \equiv \frac{\sigma}{\sigma - 1} S_{on} + \sum_{f \neq o} S_{fn} \mu_{fn}. \tag{30}$$

Nocke and Schutz (2018a) show in propositions 3 and 4 that, for demand in a class that includes our nested CES, the consumer surplus distortion from oligopoly is linear in the Herfindahl.

Herfindahl index to estimate $1/\sigma - 1/\eta = -0.444$ as the coefficient in a regression of sectoral markups on sectoral concentration.