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

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

November 12, 2019

Still preliminary & in progress

Abstract

Accumulating evidence points to rising market concentration and higher markups but has not produced a consensus on the amount of harm caused to consumers. We apply recent methods from international economics that can be implemented over many products and markets because of their minimal data requirements. The power of these methods is illustrated using two product categories where rising concentration is a plausible concern: beer and spirits. We estimate that when foreign firms take over local brands, they tend to raise costs and lower appeal. Using the estimated model, we simulate the consequences of counterfactual national merger regulation. The beer price index would be 4–7% higher had the USDOJ not forced divestitures. On the other hand, 14–30% savings could have been obtained in South America by emulating the pro-competition policies of the US and EU.

*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, and the Western Economic Association, for many helpful comments and suggestions.

[†]Sauder School of Business, University of British Columbia, vanessa.alviarez@sauder.ubc.ca

[‡]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 striking recent findings of rising concentration and markups. Grullon et al. (2019) report that indexes of concentration increased in three quarters of US industries from 1997 to 2014. De Loecker and Eeckhout (2017) and De Loecker and Eeckhout (2018) show rises in average markups in the US (from 1.2 to 1.7) and globally (from 1.1 to 1.8) since 1980. These 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 seems natural to consider a role for globalization. Conventional wisdom dictates that increases in markups must have occurred *despite* globalization, not because of it. The presumption is that increases in competition, due to reduced trade and investment frictions, have markup-shrinking effects.[1]

We can think of three arguments that challenge the standard notion that globalization should suppress markups. First, as much trade takes the form of imported inputs, decreases in the costs of such trade will tend to lower overall costs of production. When firms fail to pass on completely those cost reductions, markups rise.[2] A second argument comes from considering more carefully compositional effects that could underlie rising *aggregate* markups. As frictions separating competitors fall, "superstar" firms make inroads in the markets served by the weaker firms. Increases in market shares for the firms that already charge higher markups, leads to higher aggregate markups.[3] A third argument focuses on the way globalization has fostered more cross-border acquisitions. As large multinationals absorb previously competing entities, the acquiring firms have the ability and the incentive to increase markups.

This paper focuses on the potentially anti-competitive impacts of multinational mergers and acquisitions (M&A). We also consider the second argument, involving superstar effects, which we define as organic increases in the market shares of the largest firms. Two examples from the beverage industry illustrate how the two arguments could be confounded. The world's largest beer brewer, AB InBev, and the largest spirits distiller, Diageo, both expanded impressively over the last decade. Since both theory and evidence

[1] Brander and Krugman (1983) is a pioneering model of the "pro-competitive" effects of trade liberalization, analyzed under broader sets of assumptions in Zhelobodko et al. (2012) and Arkolakis et al. (2017).

[2] De Loecker et al. (2016) show that Indian tariff reductions led to rising markups through this channel and the 2020 *World Development Report* finds that global value chain (GVC) participation has increased markups of large corporations in developed countries.

[3] Autor et al. (2017) propose the "superstar firm framework" and marshal evidence supporting it. (Syverson, 2019a, p. 27) , and (Berry et al., 2019, p. 58), develop variations on the argument.

imply they had high markups to start out with, aggregate markups would be expected to increase. While observed markup increases align with the predictions of the superstar story, we argue that the main impact on aggregate markups in the beer and spirits markets has not been a rise in market shares of superior products, but rather the consolidation of brand ownership.

A key to understanding the market power effect of international mergers is found in the market interactions between brands referred to as "global giants" and "local stars." The former comprise the brands the MNCs sells across many markets 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, would have had no direct impact on the optimal markup had Diageo not already been selling the global giant Johnnie Walker there. The combination of the two brands' market shares is what motivates Diageo to elevate and harmonize their markups. Counterfactuals run within our oligopoly model repeatedly point to large price increases when beer and spirits multinationals bring global giants and local stars under unified ownership.

The central exercise in this paper is to compute markups under the observed set of ownership relationships and in a counterfactual where we reset every brand's owner to firm that owned it a decade earlier. 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 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 and estimated production function parameters. We follow the first approach here for three reasons. First, we lack data on input use that is critical for the production function approach. Second, a fundamental object of interest is the different markups firms charge in different markets, something the production function cannot provide since input use data are almost never linked to output markets. Third, the structure imposed in the first method is well-suited to computing counterfactual outcomes.

The precise model we use combines elements from Atkeson and Burstein (2008); Edmond et al. (2015); Hottman et al. (2016); Nocke and Schutz (2018b). The key features are multi-product oligopoly and nested constant elasticity of substitution (CES) demand. While the IO literature mainly uses random coefficient logit demand, we believe the nested CES has advantages of high tractability and low data requirements. These features permit us to replicate the analysis across 75 national markets.[4]

[4]The CES model imposes stronger restrictions on substitution elasticities than the random coefficients

Our paper contributes three key findings to the literature on the consequences of M&A. First, by considering 75 different countries, each of which has different initial market structures, we show that the same set of mergers will have highly heterogeneous impacts on consumers. The largest model-implied increases in the price index, 10–40%, occur in a handful of countries characterized by high initial concentration and passive competition agencies. Second, superstar effects do not play a large role in either beer or spirits markets over this period, being roughly balanced by convergence effects. Third, we find little evidence that new owners improve the quality or cost performance of the brands they acquire. This result echoes the finding of Blonigen and Pierce (2016) although the methodologies are entirely different. They study productivity and markups before and after acquisitions, whereas we scrutinize the firm fixed effects in regressions explaining brand performance.[5] While the specific identity of a brand's owner appears to have little impact on its inferred cost-adjusted quality, there is a systematic cost increase (statistically significant in the case of spirits) incurred by moving the location of headquarters to a country other than the brand's origin.

In addition to the substantive findings described above, our paper makes two methodological advances. Most importantly, we show how to adapt the Exact Hat Algebra (EHA) approach pioneered in Dekle et al. (2008) to conduct 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 points to a framework for addressing oligopoly issues that is more economical in its data requirements than the traditional industrial organization approach. The other method contribution is a simple way to estimate the upper level elasticity of the now commonly employed 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 accounting data on markups.

The data we work with helps us to fill some gaps in the existing literature on concentration and markups. 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-

methods preferred in the IO literature. However, Head and Mayer (2018) show that a CES model (calibrated to replicate the observed average elasticity of substitution between brands) does a good job of approximating aggregate outcomes of rich substitution models in counterfactual simulations.

[5] Our results and approach are also consistent with the major synthesis of retrospective merger studies conducted by Kwoka (2014).

industry, leather tanning should be thought of as an input to the other two. We focus on consumer beverages where nearly all varieties relate to each other as substitutes. Indeed IO economists have used these very markets in their studies, contributing estimates of the own price elasticities that we use to calibrate our model. Our dataset is also well equipped to measure market concentration as it overcomes two main limitations. First, instead of using firms' output—which includes exports and excludes imports—we calculate measures of concentration from brand-level sales in a given market without regard to where the goods were sourced. Second, we depart from the focus on the U.S., and measure increases in aggregate markups across 75 countries.

The remainder of the paper proceeds as follows. Section 2 describes the data, highlighting the major patterns of worldwide expansion of the largest beer and spirits makers. We show how oligopoly Lerner indexes vary with conduct under nested CES demand in section 3. There we also show how we back out brand-market appeal, cost-adjusted appeal, and markups. Section 4.2 estimates the effects firm ownership effects on brand performance. Armed with an estimation of how changes in corporate headquarters of a brand affect perceived quality and costs, section 5 computes a counterfactual that reverses the last decade of transactions in brand ownership, restoring the 2007 owners of every brand in every market. This allows us to quantify the impact of multinational brand amalgamation on consumer welfare and aggregate markups.

2 Data: sources and patterns

Our data set combines four distinct components. The first data set provides sales at the brand-market-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 headquarter country for each brand-owning firm. Finally, we use standard data (available from CEPII) on bilateral proximity measures.

2.1 Market shares and prices

Passport Global Market Information Dataset (GMID), from Euromonitor, records sales information for 83,000 brands owned by 46,000 companies, across 153 product categories, in 79 countries for 10 years. 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." The latter often accounts for a sizeable share of total sales as much as 31% for beer in Germany. We calculate firm market shares S_{fn} using as a denominator the sales of all brands—including others and private label.

GMID also allows us to track any changes in ownership, at the brand level, occurring over the period 2007–2018, which is a unique feature of this dataset. 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.

This dataset is better suited for measuring market concentration than data from the economic census or firm-level databases such as Compustat and Orbis. This is because it overcomes two main limitations inherent to databases relying on firm's revenue. First, firm's revenue includes exports to other markets and excludes imports. Thus it does not measure sales in the market in question.[6] To the extent that imports comprise products of foreign firms, this will lower concentration in the market; but, imports carried out by large firms with little or no local production can actually increase concentration relative to measures based on domestic shipments. Similarly, including exports could significantly upward bias concentration when the exporters are multinationals that use the local market as export platform. Our data overcomes these issues, as we construct measures of concentration by looking at brand-level sales in a market without regard to where the goods being sold were sourced.[7] Studies of concentration using Compustat omit private companies, which include a few large firms (e.g. Bacardi) and the often large fringes of small firms. Both Compustat and census miss sales of multi-category companies outside their assigned SIC.[8] Third, most of the research on rising concentration only uses data pertaining to the US market, whereas our analysis comprises the 75 countries for which quantity and value data are available for both spirits and beer.

The full GMID data set has too many brands and companies for it to be practical to identify brand origins and firm headquarters. We therefore reduced the products under consideration to seven beverages shown in Table 1. This set of goods have the added attraction of being the only ones in GMID with both sales values and sales volumes. By combining the two series, we can back out the demand shifter for each brand. Table 1 shows that each category comprises hundreds of firms and there are thousands of brands in most categories.

The regression method we use to estimate firm ownership effects on brand perfor-

[6]Compustat has the even 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.

[7]We have no systematic data on the location of production.

[8]Compustat classifies Pernod Ricard, the 2nd 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	multi-n	multi-f		HQ	Origin	Market
Beer	2486	383	807	510	80	95	78
Spirits	3280	626	597	958	90	108	77
Wine	1840	254	277	857	57	56	54
Water	1404	236	275	906	84	98	92
Carbonates	1107	267	214	509	84	88	92
Coffee	868	220	228	570	82	83	92
Juice	1440	355	319	942	89	95	92

mance is only identified when brands are "mobile" across firms and markets. Table 1 shows that brand ownership changes are common, with the highest count and fraction being seen in the beer industry. 32% of the beer brands in the data set had more than one owner. This includes some cases of brands, such as Corona and Fosters, that have different owners in different markets. Spirits also exhibits substantial mobility of brands across owners, with about 18% having more than one owner. Spirits has the highest count of multi-market brands, which is important for backing out both brand effects and brand-origin frictions. For all these reasons, our estimation will focus on beer and spirits, though we report results for other beverages in the appendix.

Figure 1: Global giants are rare

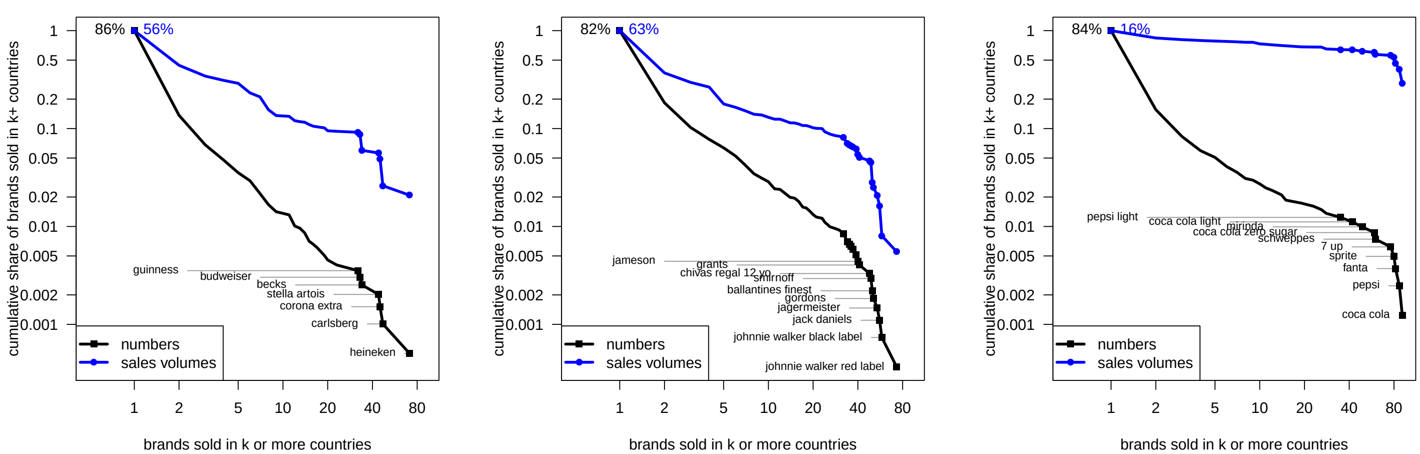

(a) Beer (b) Spirits (c) Carbonates

Note: Symbols mark brands sold in $>$ 30 countries. They account for 9%, 8%, and 64% of volume in Beer, Spirits, and Carbonates. Log scales on both axes. Calculations exclude Others (counts, destinations not known).

Figure 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.[9] The interesting thing is that Beer and Spirits exhibit *much* less of this global brand concentration than Carbonates. For all three products, brands selling in more than 30 markets are rare—0.4% of Beer, 0.9% of Spirits, 1.2% of Carbonates—but in Carbonates those few brands make 64% of all sales, whereas in Beer and Spirits, it is under 10%. The flip side of this is that single-market brands, which constitute over 80% of brands for all three goods, are relatively unimportant in world sales of Carbonates (16%) whereas they account for a substantial majority of Beer and Spirits sales. While most single-market brands have low market shares, a few—the local stars—are the leading brands in most markets. This distributional pattern of sales means that the potential for increases in market power via multinationals buying local brands is much bigger in Beer and Spirits than in Carbonates.

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 headquarter country of each company in GMID dataset is obtained by combining Orbis from Bureau van Dijk, the historical Directory of Corporate Affiliations from Lexis-Nexis, and Capital IQ. These datasets provide detailed information about ownership structure of the firm, as well as information on their affiliates' location, sectors of operation, sales, assets, operating profits and employment. Matching the name of each brand's owners in the GMID dataset with the names of firms in the Orbis, Lexis Nexis, and Capital IQ datasets, we take the headquarter 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.

The origin of a brand is the country where it was developed and introduced. Thus Labatt is a Canadian brand and Corona Extra is a Mexican brand even though both are currently owned by AB InBev NV, a Belgium headquartered owner. 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,

[9] 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.

[10] The 49% of China Resources owned by SABMiller is insufficient to show SABMiller as the owner of the Snow brand in China (before it was forced to divest its share).

the categories with the most brands, we often relied crowd-sources websites that rate each product.

2.3 Visualizing multinational brand amalgamation

Table 2 provides a preview of the type of information in our data and also motivates the title of our paper. Diageo was formed in 1997 as a merger of Grand Metropolitan and Guinness. It dramatically expanded its portfolio of spirits brands when took over the brands of the failing Seagram company in 2001. On its website Diageo distinguishes between "Global Giants" and "Local Stars." We operationalize the concept of the former as brands that are sold in many countries and achieve high market share world wide and the latter as brands sold in few markets but which achieve very high market share in their country of origin. While the table shows all of Diageo's Global Giants, it selects 7 of the Local Stars to illustrate the range of countries represented. In keeping with Diageo's own focus, 12 of the 14 brands shown are spirits.

Table 2: A selection of Diageo brands

Global Giants

	Johnnie Walker	Gordon's	Tanqueray	Smirnoff	Captain Morgan	Baileys	Guinness
Origin:	UK	UK	UK	Russia	Jamaica	Ireland	Ireland
# Markets:	68	21	28	64	43	57	30
rank (world):	2nd	30th	46th	1st	12th	24th	21st
born (bought):	1860 (1997)	1769 (1997)	1830 (1997)	1864 (1987)	1944 (2001)	1973 (n/a)	1759 (1997)

Local Stars

	Ypióca	McDowell's No1	Yeni Raki	Cacique 500	Bundaberg Rum	Crown Royal	Tusker Lager
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)

Note: 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.

A striking aspect of the brands shown in Table 2 is they are mainly very old and only one was invented by its current owner. The exception is Bailey's Irish Cream which was invented in 1973 within a division of Grand Metropolitan, which went on to become Diageo after the 1997 merger with Guinness. The other brands in the table originated from 56 to 260 years ago. For Diageo the main way to expand its brand portfolio has been the acquisition of brands invented by other firms.

Figure 2: The growth of beer multinationals

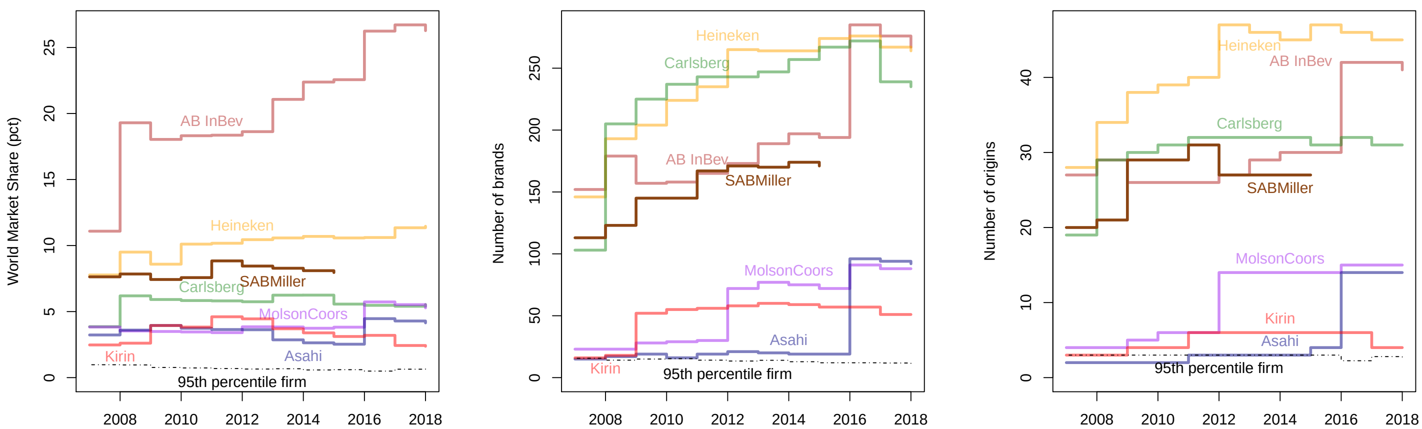

Note: In 2008 InBev purchases Anheuser-Busch and Heineken and Carlsberg jointly purchase Scottish & Newcastle (along with BBH) and redistribute the acquired brands among themselves. In 2009 AB InBev sells off Korean and East European brands (forming Starbev) and Kirin acquires Lion (NZ). In 2012 MolsonCoors 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 and 3 illustrate the rise in market shares, brand ownership, and diversity of brand origins for the seven largest companies in the beer and spirits industries.

3 CES oligopoly model

The model we use is guided by the data described above. There are a finite number of firms who compete oligopolistically and have observed market shares for each of the brands they sell. The next two subsections describe the assumptions we make about demand and market conduct.

3.1 Demand

Consumers have CES preferences over brands offered in country n, and they substitute across brands with an elasticity σ, regardless of the identity of the firm that owns the brand. This differs from the two-nested CES structure Hottman et al. (2016) which is Cobb-Douglas between sectors but has σ^F as the firm-level CES and and σ^U as the bar-

Figure 3: The growth of spirits multinationals

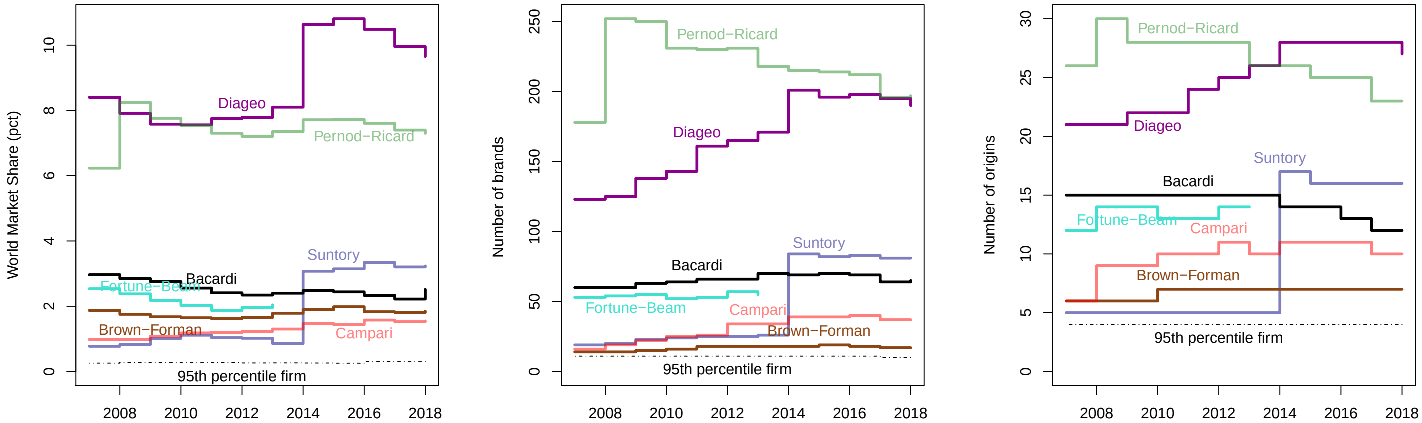

Note: In 2008 Pernod-Ricard buys Vin & Sprit (owner of Absolut and 74 other brands). In 2014 Suntory buys Beam (which had been spun off from Fortune Brands in 2011) and Diageo buys UB Group.

code level CES. Their estimation results show that σ_g^F is systematically lower than σ_g^U.[11] If brands never changed owners, this extra layer of nesting could be added without changing markups since the latter only depend on the firm-level market share.[12] Changes in brand owner are the focus of our paper. We do not find it reasonable that a brand would change its σ with respect to other brands merely because of change in ownership. To be concrete, Smirnoff does not become a closer substitute for Tanqueray simply by changing ownership from Seagram to Diageo (the owner of the latter brand).

The appeal of a brand, A_{bnt}, is market, n, and time, t, specific. First, having a brand appeal that is market-dependent allows the model to capture the fact that a brand can be popular in one country (very often its origin), but be less attractive to consumers in other countries. In section 4.2 we estimate how much of the variation in brand's appeal across countries can be explained by a brand being particularly appealing to consumers in the same country as where the brand was originally designed, this is we estimate the importance of "home bias." Second, having a time-variant appeal allow us to capture changes in taste, if any, that consumers experience when a brand changes owners after an acquisition.[13] For simplicity, we suppress the time subscripts for the rest of this section.

Let the representative consumer in market n have a utility taking a nested Constant

[11]One explanation is that bar-codes frequently differ just because of packaging in different volumes which one think makes two bar-codes of the same firm closer substitutes than bar codes from different firms.

[12]Nocke and Schutz (2018b) have an appendix where they explore more generally whether intra-firm nests matter.

[13]Consumers could change their valuations of the brand because they perceived an actual change in the quality of the product after the acquisition, or simply because their valuation of the brand also factors in their perception of its owner.

Elasticity of Substitution (CES) form. At the upper level, consumers allocate their income among 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 n as

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

where P_{gn} is the price index of sector g in market n, P_n is the overall price index, and X_n is aggregate expenditure.

At the lower level, inside g, the quantity index Q_{gn} is given by

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

where q_{bn} denotes the quantity consumed of each brand b in market n. Each brand is implicitly associated with a unique sector g, so we can dispense with g subscripts on all variables that are indexed by b.

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

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

where p_{bn} is the price of brand b in market n, and P_n is the market price index which aggregates over all the brands offered in the market n, as indicated by $\mathbb{I}_{kn}$:

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

The total market share of firm f in market n, S_{fn}, is obtained by aggregating the market shares of all the brands the firm's portfolio ($\mathcal{F}_f$) and offers in market n ($\mathbb{I}_{bn} = 1$):

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

As shown in table 1, in the cross section the extensive margin of where brands are offered is very important. However over the decade of data we have, there is not much action across time in $\mathbb{I}_{bnt}$. Appendix section A documents the very low rates of adding and dropping brands across markets. Since it does not appear to be an important aspect of the data and would prevent us from using exact hat algebra for the counterfactuals, we treat $\mathbb{I}_{bnt}$ as an exogenous characteristic of brands, like their appeal and production cost.

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 marginal cost of delivering brands to 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 owns:

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

3.2 Markups for different conduct assumptions

Regarding the pricing strategy of firms, we follow 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, but treat the aggregate expenditure and price level (X_n and P_n) as given.

It is 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}:

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

The results that markups are the same across all the brands a firm owns under CES demand is pointed out in Hottman et al. (2016) and contrasts sharply with the case of linear demand analyzed by Mayer et al. (2014).

The functional form of ϵ_{fn} (and hence μ and L) depends on the assumed mode of oligopoly conduct. The Lerner indices implied by the two standard conduct assumptions are

$$\underbrace{L_{fn} = \frac{1}{\sigma - (\sigma - \eta)S_{fn}}}_{\text{Bertrand}} \quad \text{and} \quad \underbrace{L_{fn} = \frac{1}{\sigma} - \left(\frac{1}{\sigma} - \frac{1}{\eta}\right) S_{fn}}_{\text{Cournot}} \tag{10}$$

In addition to the firms selling listed brands, the GMID data contain a sometimes large share of sales attributable to other brands, whose owners are not reported. Since the threshold below which GMID aggregates brand sales into a group called "Others" is so low (0.1%), we treat those brands as a part of a monopolistically competitive fringe. Thus, for other brands, $L_{0n} = 1/\sigma$.

Prices can be expressed in terms of either markup:

$$p_{bn} = \mu_{fn} c_{bn} = c_{bn}/(1 - L_{fn}) \tag{11}$$

Although the CES oligopoly model lacks closed-form solutions for prices, equilibrium can be obtained via fixed point iteration, starting with a guess of prices (such as the monopolistic competition price vector $p^0_{bn} = (\sigma/(\sigma-1))c_{bn}$). At each step, market shares are obtained for a given set of markups, which then imply a new set of optimal markups, new market shares, until convergence to unique price and market share vectors is reached.

A major attraction of the CES oligopoly model is that it provides simple expressions for the 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

Fortunately, for the purpose of this paper, industrial organization economists have already devoted considerable efforts to the estimation of brand-level own-price elasticities for the very products we study. We will treat those estimated elasticities as moments used to pin down σ for each of the categories we consider.

Table 3: Estimates of own-price elasticities and implied elasticities of substitution

Product group	Mean σ	Mean ϵ_b	# Estimates	# Papers
Beer	4.49	4.48	9	5
Spirits	3.38	3.37	9	2
Wine	5.18	5.17	1	1
Carbonates	3.44	3.43	6	5
Coffee	3.96	3.92	5	4
Juice	3.88	3.85	2	2
Water	2.66	2.65	3	1

The IO literature summarized in Table 3 reports mean or median of brand-level own price elasticities, estimated from the demand side of the model before imposing any market structure. Those demand elasticities cannot be interpreted as direct estimates of the elasticity of substitution σ_g because of non-negligible market shares. We can, however, use the brand-level formula for CES own-price elasticity $\epsilon_b = \sigma_g - (\sigma_g - \eta)s_b$ and invert it to solve for σ_g as a function of either the mean or the median (denoted with function $m_g()$ of estimated demand elasticities in the category:

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

We hold these σ_g constant over time.

Atkeson and Burstein (2008), the pioneering work using nested CES oligopoly, impose $\eta = 1.01$ and consider $\eta = 1.5$ in a sensitivity analysis. Because this parameter is so important in our quantification of the markups, we want to discipline it with data from the industries we study. In contrast to the abundance of brand-level elasticities estimates, we could not find appropriate cross-category elasticity estimates in the literature. However, the Bertrand and Cournot markup formulae in equation (10) can both be inverted to obtain η as a function of σ_g, S_{fn} and L_{fn}. Since we have σ_g and data on S_{fn}, the last piece of data needed to retrieve an implied η would be L_{fn}, the firm-market Lerner index. Unfortunately, accounting data are generally unavailable at the market level because firms report their "consolidated" accounts, aggregating over all markets they serve. We therefore aggregate national Lerner indices to the firm level, denoted L_f. At the firm level, the consolidated markup is the ratio of worldwide profits over worldwide sales:

$$L_f = \frac{\sum_n S_{fn} L_{fn} X_{gn}}{\sum_n S_{fn} X_{gn}} \tag{12}$$

For known σ_g, firm-level market shares, and market-level expenditures (X_{gn}), all firm-level Lerner indexes (L_f) are a function of the same unknown parameter η. Matching the predicted Lerner indexes with the observed Lerner indexes inferred from accounting data, we construct a loss function as the sum of squared deviations between the theoretical and accounting measures of L_f.

Compustat provides consolidated revenue and expense data for the seven largest multinationals in the beer and six of the seven largest spirits makers.[14] The accounting measure of L_f is computed as the ratio of gross profits over sales using consolidated income statements. The critical issue is which costs to subtract in calculating profits. The

[14] The third largest spirits producer, Bacardi, is privately owned and does not provide the required data.

theory dictates it should be only marginal costs but, as discussed in Basu (2019) and 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. Following De Loecker and Eeckhout (2017), we take COGS to be entirely variable costs. Our conservative markup measure treats all of SGA as variable costs as well, 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.[15] 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 (high end) markup deducting only 32% of SGA.[16]

Our η estimate minimizes the distance between the midpoint of the Bertrand and Cournot computations of L_f and the midpoint of the conservative and liberal accounting Lerner indexes.[17] We compute the theoretical and accounting L_f for each of the seven largest publicly traded multinationals in those two industries over the 2007–2018 period. There are 157 observations (some firms are absorbed via mergers, leading to an unbalanced panel). The distance-minimizing value of η is 1.62, which corresponds to a monopoly Lerner index of 62%.

With σ_g and η in hand, we can graphically compare the theoretical markups to those obtained from accounting data. Figure 4(a) graphs the Lerner index functions for Bertrand and Cournot conduct assumptions. The blue lines use our σ estimate for beer (4.5) whereas the red line uses our Spirits estimate ($\sigma = 3.4$). In both Cournot and Bertrand, L ranges from $1/\sigma$ for $S_{fn} = 0$ (the monopolistic competition benchmark) to $1/\eta$ for $S_{fn} = 1$ (monopoly). For a given product, the Lerner index for Bertrand lies under the corresponding index for Cournot for $0 < S_{fn} < 1$.

Figures 4(b) and (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 theo-

[15] 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.

[16] To clarify, we express the accounting markups in terms of the underlying Compustat variables: $L_f^a = (\text{sale} - \text{cogs} - \vartheta\text{xsga})/\text{sale}$. Let $\vartheta = 1$ correspond to the conservative markups and $\vartheta = 0.32$ to the liberal markups. In the few instances in Compustat where xsga is incomplete, we replace it with operating expenses (xopr) minus cogs.

[17] This avoids the awkwardness of estimating four different values of η corresponding to each of the conduct-accounting combinations.

Figure 4: Oligopoly markups for Bertrand and Cournot, compared to accounting data

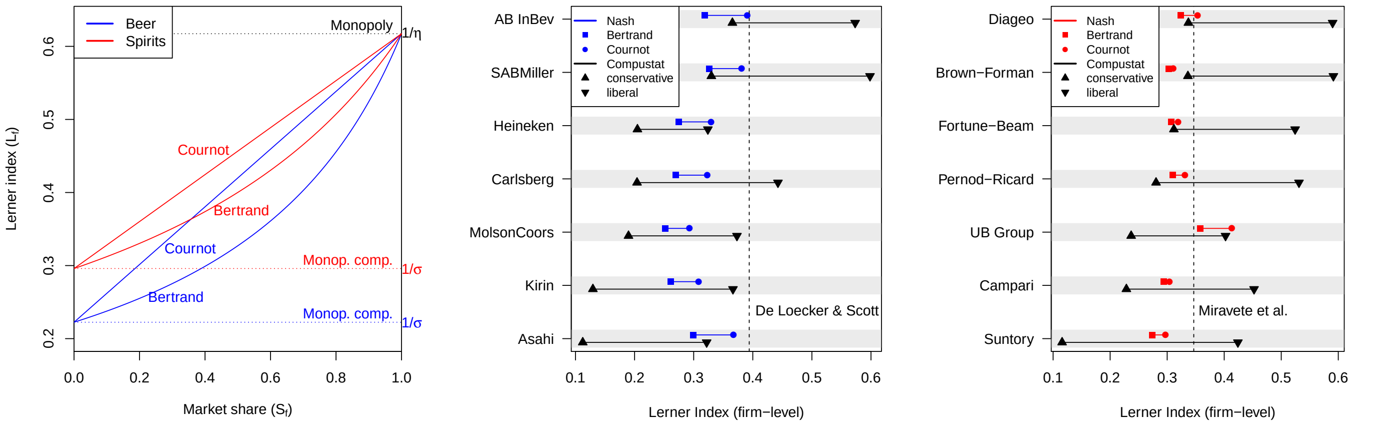

(a) Theory (b) Beer data (c) Spirits data

retical interval, we show the range between our lower and upper bounds for accounting markups (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.[18]

There are three salient points in the markup figures. The accounting and theory intervals overlap for every beer maker and for all but one (Brown-Forman) spirits maker. For beer, AB InBev and SABMiller have both the highest predicted markups (based on market shares) and the largest accounting markups. The fact that the theoretical markups (based on calibrated σ_g and η) are broadly consistent with the accounting data provides evidence 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.

[18]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) ranging 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 $L = 1 - 1/1.65 = 0.39$.

3.4 Backing out brand type and appeal

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.[19] 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. 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 unlike the basic Melitz model, the data reveal large variation in φ_{bn} across markets.

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

$$\varphi_{bn} = s_{bn}^{1/(\sigma-1)} \mu_{f(b)n} P_{gn}. \tag{13}$$

In order to isolate brand type as a function of observables, we need to eliminate the price index, which can be accomplished by dividing by any other $\varphi_{b'n}$ (or index of brand types) since they all depend on the same price index.

Aggregating the brands in each $\mathcal{F}_f$ portfolio that firm f offers in market n,

$$S_{fn} = \sum_{b \in \mathcal{F}_f} \mathbb{I}_{bn} s_{bn} = \mu_{fn}^{1-\sigma} P_n^{\sigma-1} \sum_{b \in \mathcal{F}_f} \mathbb{I}_{bn} \varphi_{bn}^{\sigma-1} = (1 - L_{fn})^{\sigma-1} T_{fn} P_n^{\sigma-1}, \tag{14}$$

where T_{fn} is the firm-market level aggregator of brand characteristics, called firm type by Nocke and Schutz (2018b) since they only consider a single market, whereas here firm type varies by market.

$$T_{fn} = \sum_{b \in \mathcal{F}_f} \mathbb{I}_{bn} \varphi_{bn}^{\sigma-1} \tag{15}$$

The key point about T_{fn} is that it is a sufficient statistic for the performance (market share, profit) of the firm. Computationally, this means that equilibrium firm market shares can be calculated without considering individual brands if T_{fn} is known.

The market share of other brands, a monopolistically competitive fringe, is $S_{0n} = \mu_{0n}^{1-\sigma} T_{0n} P_n^{\sigma-1}$. Inverting, we have $T_{0n} = \mu_{0n}^{\sigma-1} S_{0n} P_n^{1-\sigma}$. Therefore, we can normalize the

[19] Melitz (2003) made this point 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.

measure of cost-adjusted quality φ_{bn} in equation (13) by $T_0^{1/(\sigma-1)}$ to obtain

$$\breve{\varphi}_{bn} = \frac{\varphi_{bn}}{T_{0n}^{1/(\sigma-1)}} = \left(\frac{s_{bn}}{S_{0n}}\right)^{1/(\sigma-1)} \frac{\mu_{f(b)n}}{\mu_0}, \tag{16}$$

where $\mu_0 = \sigma/(\sigma-1)$ in all markets. Markups for all other brands are obtained by applying a conduct assumption inside equation (10, and using $\mu_{fn} = 1/(1-L_{fn})$. $T_0^{1/(\sigma-1)}$ is a CES index of the φ of the unlisted brands. If there were a single other brand, indexed 0, it would have $T_0^{1/(\sigma-1)} = \varphi_0$.

With data on brand prices as well as market shares, one can further separate out brand appeal (A_{bn}). To see this, take logs of equation (4) yielding

$$\ln s_{bnt} = (\sigma-1)\ln A_{bnt} - (\sigma-1)\ln p_{bnt} + (\sigma-1)\ln P_{gnt}. \tag{17}$$

Since the price index is a gnt variable, it is common to all brands in a given product-market-year and can therefore removed through demeaning. As in Hottman et al. (2016), a tilde over a variable denotes its geometric mean over the relevant market-year (specified in its subscript).[20] So long as we have an estimate of σ we can express inferred appeal as a function of observables:

$$\ln(A_{bnt}/\tilde{A}_{gnt}) = \frac{\ln(s_{bnt}/\tilde{s}_{gnt})}{\sigma-1} + \ln(p_{bn}/\tilde{p}_{gnt}). \tag{18}$$

Only relative A_{bnt} within a product-market-year can be identified since multiplying all the A_{bnt} by a scalar would not change the equilibrium market shares conditional on prices.

Equation (18) is equivalent to the regression approach of Khandelwal et al. (2013) equation (7) except that they aggregate over multiple sectors (and therefore include sector fixed effects), whereas we calculate appeal within each category of goods. Equation (18) is also equivalent to a logged version of Redding and Weinstein (2018) equation (17).

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 the firm improves performance everywhere, and a localized effect that depends on the proximity of the firm's HQ to each market served by the brand. To iso-

[20] When calculating the geometric means of market shares and prices, we include only the individually "listed" brands.

late 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. To isolate those components of brand performance , we estimate high-dimensional fixed effects regressions familiar in the labor literature initiated by Abowd et al. (1999).

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:

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

Substituting for price in equation (17) and applying the definition of brand type, we have

$$\ln s_{bnt} = (\sigma - 1)\left[\ln \varphi_{bnt} - \ln \mu_{f(b,t)nt}\right] + (\sigma - 1)\ln P_{gnt}. \tag{19}$$

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

$$\ln \varphi_{bnt} = \ln \varphi_b^B + \ln \varphi_{f(b,t)}^F + \ln \delta_{i(b)n}^B + \ln \delta_{h(f(b,t))n}^F + \varepsilon_{bnt}. \tag{20}$$

The function $i(b)$ maps brand b to its origin i, which is time invariant. On the other hand, ownership changes over time, implying that the mapping of a brand to the headquarter country of its owner, $h(f(b,t))$ depends on time.[21]

The δ^B and δ^F capture the impact of observable frictions on φ_{bnt}. We have in mind 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 estimate the magnitudes of such effects using two "home" variables. The first home$_{in}$ takes a value of 1 when brands sold

[21] Ownership of a brand sometimes differs across markets, for example when competition authorities force divestitures. We omit this infrequent case in the notation but take it into account in the estimation and counterfactuals.

in their country of origin ($i = n$). The second home variable is defined at the headquarter level and equals one when the owner of the brand has its HQ in the market ($h = n$). We also include common language and the log of distance, with both defined in terms of in and hn.

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 the brand owner's headquarters.

The final estimating equation for cost-adjusted appeal uses our inferred values, $\check{\varphi}_{bnt}$ from (16) in place of the unobservable φ_{bnt}.

$$\ln \check{\varphi}_{bnt} = \mathrm{VFE}_b^B + \mathrm{VFE}_{f(b,t)}^F + \mathrm{VFE}_{gnt} + \mathbf{X}'_{i(b)n}\mathbf{d}^B + \mathbf{X}'_{h(f(b,t))n}\mathbf{d}^F + \varepsilon_{bnt}, \tag{21}$$

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 VFE (varphi fixed effect) have structural interpretations as $\ln \varphi_b^B$, $\ln \varphi_{f(b,t)}^F$, and $-\ln T_{0gnt}/(\sigma - 1)$. To determine the effect of each friction variable working through the demand side alone, we also estimate a version of equation 21 where $\ln A_{bnt}$ replaces $\ln \check{\varphi}_{bnt}$ as the dependent variable. The differences between the coefficients in those two regressions corresponds to the cost effect.

The primitive determinant of brand market shares in equation (19) 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

$$\ln \mu_{bnt} = \mathrm{MFE}_b^B + \mathrm{MFE}_{f(b,t)}^F + \mathrm{MFE}_{gnt} + \mathbf{X}'_{i(b)n}\mathbf{g}^B + \mathbf{X}'_{h(f(b,t))n}\mathbf{g}^F + \upsilon_{bnt}. \tag{22}$$

In this regression, the coefficients do not reveal structural parameters because of the non-linear mapping from frictions to 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 19, we have the estimable log market share equation:

$$\ln s_{bnt} = \text{SFE}^B_b + \text{SFE}^F_{f(b,t)} + \text{SFE}_{gnt} + \mathbf{X}'_{i(b)n}\mathbf{r}^B + \mathbf{X}'_{h(f(b,t))n}\mathbf{r}^F + \xi_{bnt}. \tag{23}$$

The additive-in-logs structure implies that market share friction coefficients are algebraically tied to the $\ln \check{\varphi}_{bnt}$ and $\ln \mu_{bnt}$ coefficients via $\mathbf{r} = (\sigma - 1)(\mathbf{d} - \mathbf{g})$. Similarly, the coefficients on $\ln \check{\varphi}_{bnt}$ and $\ln \mu_{bnt}$ for different conduct assumptions are linked through equation (16): 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. We will see that the R^2 of the market share equation is substantially below one, mostly explained by inability of the model to fully explain variation in appeal. This is not surprising for any traveller who has noticed certain brands are inexplicably popular in certain countries. The fact motivates the usefulness of exact hat algebra for counterfactuals since this method implicitly takes into account the unobserved determinants of market share that are invariant to the counterfactual.

4.2 Estimation results

Table 4 reports results for regressions that pool Beer and Spirits brands. The most striking result is that, on average, home-origin brands have huge advantages. Since $\exp(1.039) \approx 2.83$, home increases market share by 183%. The largest impact comes on the taste side (home bias). In particular being a home brand raises demand equivalent to a 25% price change. Brands from faraway countries also lower cost-adjusted appeal, with a distance elasticity of -0.045. The corresponding elasticity for cars in Head and Mayer (2019) is -0.088. Brands also have significantly higher market shares in their HQ country.

The pooled regressions in Table 4 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 5 graphs the results, expressed as *ad-valorem* equivalents of the home advantage for cost-adjusted appeal (φ).[22] The home bias estimated under the Cournot conduct assumption is systematically larger

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

Table 4: 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	1.039^{a}	0.223^{a}	0.355^{a}	0.016^{a}	0.370^{a}	0.031^{a}
	(0.133)	(0.073)	(0.050)	(0.004)	(0.052)	(0.006)
distance	-0.124^{a}	0.035	-0.045^{a}	-0.002^{b}	-0.047^{a}	-0.004^{b}
	(0.037)	(0.022)	(0.015)	(0.001)	(0.015)	(0.002)
common language	0.051	−0.054	0.009	−0.0001	0.009	0.0002
	(0.078)	(0.050)	(0.031)	(0.002)	(0.032)	(0.003)
home (HQ)	0.342^{a}	0.102^{c}	0.178^{a}	0.032^{a}	0.202^{a}	0.056^{a}
	(0.110)	(0.062)	(0.043)	(0.004)	(0.045)	(0.007)
distance (HQ)	0.026	0.011	0.016	0.001	0.015	0.001
	(0.033)	(0.020)	(0.013)	(0.001)	(0.013)	(0.001)
com. lang. (HQ)	0.102^{c}	0.049	0.050^{b}	0.003	0.053^{b}	0.006
	(0.061)	(0.039)	(0.025)	(0.003)	(0.026)	(0.004)
Observations	95,245	95,245	95,245	95,245	95,245	95,245
R^2	0.658	0.654	0.596	0.900	0.604	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 headquarter country. Significance levels: 1% (a), 5% (b), and 10% (c).

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 *ad-valorem* equivalent of the part of home bias that comes from the demand side. These appeal effects do not depend on conduct, since they are extracted directly as demand shifters.

As seen in panel (a) of Figure 5, 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 achieve high market shares without purchasing those local stars.

For beer brewers, about one third of the origin home advantage (a 25% AVE) comes from the consumer preference for domestic brands. The AVE if 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

Figure 5: Home bias by type and category over time

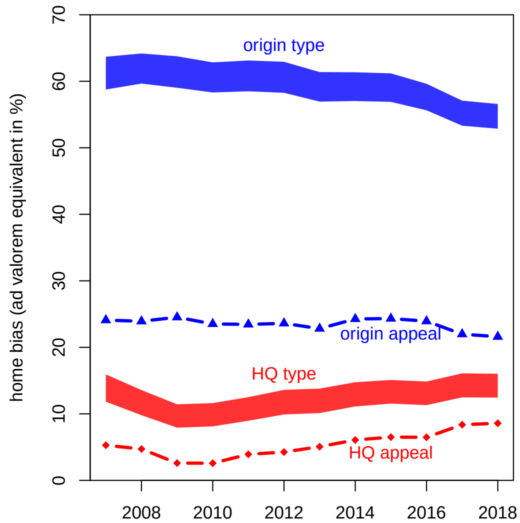

(a) Beer

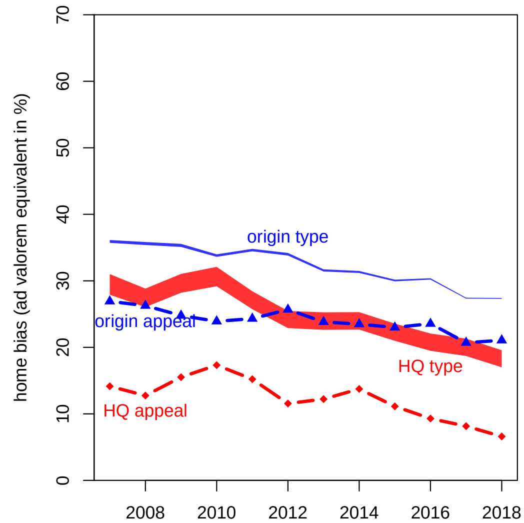

(b) Spirits

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

also produced locally, transport costs incurred by foreign brands should matter more for Beer.[23] While other papers have estimated home bias, notably for the auto industry (Coşar et al. (2018), Head and Mayer (2019)), we believe this is the first paper to quantify what fraction of the home bias is accounted for by tastes.

The other novelty is that we can estimate the home bias related to the HQ country of the brand's owner. This is important in the context of industries with large waves of brand amalgamation by foreign firms. This HQ-related home bias is estimated as equivalent to around 10–15% home advantage for Beer, and 20%–30% in Spirits. This is the immediate cost imposed on a brand when bought by a foreign company. In order to rationalize those purchases, there must therefore be some gains in the form of either synergies or increased market power.

The way we quantify brand-firm synergies is by estimating firm-level fixed effects. 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.[24] Therefore, a necessary condition for the realization of positive synergies is that there is enough variance in the estimated firm-level fixed effects. Variance is not a sufficient condition since, in addition, brands should flow from the weak to the stronger firms. Because brand

[23]This explanation is further supported by the distance effects reported in Appendix Tables 10 and 11, where the coefficient for Beer is more than twice as large as the one for Spirits.

[24]The structural interpretation of VFE$_f$ in equation (21) is $\ln \varphi^F_{f(b,t)}$. Therefore, following the transfer of b to firm A from firm B, cost-adjusted appeal rises by $\ln \varphi^F_{A(b,t+1)} - \ln \varphi^F_{B(b,t)}$.

and firm effects enter multiplicatively in the sales and profits equations, this mechanism will be stronger for the best brands, predicting positive assortative matching. In the next subsection, we measure both the variance of firm fixed effects and their association with brand effects.

4.3 Estimating the contribution of firm effects

Before relating brand and firm fixed effects, we need to establish how those can be separately identified. As is the case with firm and worker effects in the wage determination literature, identification requires "mobility." In our context, movements are changes in ownership of brands which connect different firms in the same way that workers changing jobs connect establishments in the employer-employee regressions introduced by Abowd et al. (1999).

Table 5: Brand mobility in the connected set

Product group	# Firms		Mobility		Sales share	
Beer	91	22	15.9	58.7	80.0	70.8
Spirits	94	18	8.4	33.6	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	91.3	65.7
Juice	60	2	2.7	13.0	44.5	2.8
Coffee	3	NA	2.7	NA	33.2	NA
$\geq$ 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.

Andrews et al. (2008) show that the ratio of total movements to the population of firms reveals the importance of limited-mobility bias when estimating multi-way fixed effects regressions. In the labor literature, this bias has been shown to generate patterns of negative correlations between employer and worker effects. In the third and fourth columns of Table 5, 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).

We conduct the variance and covariance analysis of fixed effects for the two products with high brand mobility, Beer and Spirits. Table 6 displays correlation between the fixed effects estimated for each brand and each firm in three different regressions using market shares, appeal and cost-adjusted appeal (assuming Bertrand) as dependent variables. The corresponding regression coefficients are shown in the first three columns of appendix tables 13 and 15. Following the approach advocated by Andrews et al. (2008) to mitigate limited-mobility bias, these regressions restrict the sample to moving brands and high mobility firms.[25] In each table, the diagonal shows the ratio of the variance of the relevant fixed effect to the variance of the dependent variable.

Firm effects explain at most five percent of the variance of performance measures for both beer and spirits. Therefore, the identity of the firm owning a brand explains very little of the variance in its market share, appeal and cost-adjusted appeal. A corroboration of this finding is that the addition of firm fixed effects to the regressions adds very little to the fit of the regressions. Appendix tables 12 and 14 present results of the same regressions as in Tables 4 and 10. The incremental R^2 for the pooled regressions with brand type as the dependent variable is just 0.008. 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 6.

The off-diagonal elements of Table 6 show the sign and magnitude of assortative matching. For beer, firm and brand fixed effects are all negatively correlated, despite adopting the recommended approach to mitigate limited mobility bias. For spirits, the correlations have mixed signs, but all have small magnitudes (under 10%). [26]

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 will not influence the price outcomes of ownership changes. Hence we do not need to take a

[25] There are no major changes in the friction coefficients compared to results not imposing that restriction, and shown in tables 10 and 11.

[26] In the appendix, Table 21 shows the equivalent of fixed effect correlation tables for regressions on the full sample. As in the labor literature, samples that do not mitigate limited mobility bias generate seemingly spurious findings of negative assortative matching.

Table 6: Correlations between fixed effects and their explanatory strength

	Brand			Firm		
Dep. var.:	share (s_{bn})	appeal (A_{bn})	type (φ_{bn})	share (s_{bn})	appeal (A_{bn})	type (φ_{bn})
Beer						
brand market share	1.172					
brand appeal	0.740	1.192				
brand type	0.987	0.732	1.259			
firm market share	-0.082	-0.064	-0.061	0.029		
firm appeal	-0.057	-0.106	-0.024	0.795	0.046	
firm type	-0.070	-0.032	-0.049	0.960	0.733	0.029
Spirits						
brand market share	0.682					
brand appeal	0.720	0.679				
brand type	0.999	0.716	0.701			
firm market share	-0.051	0.124	-0.055	0.042		
firm appeal	-0.159	-0.048	-0.172	0.693	0.045	
firm type	-0.059	0.119	-0.059	0.991	0.664	0.051

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.

stance on them in the counterfactuals when considering the consequences of mergers on the consumer surplus, an exercise to which we now turn.

5 Effects of brand acquisitions on consumers and markups

We want to quantify how much harm to competition has been caused by the process of multinational brand amalgamation. We are interested in calculating the cumulative impact on concentration and consumer surplus from the brand acquisitions by MNCs. There are many ways one could approach this. The way we follow is to reverse all the mergers that occurred over our period: The counterfactual restores 2007 owners to each brand in 2018. The simulation adjusts the market shares of all brands. In addition to taking into account how alternative ownership patterns affect firm level market shares and hence their optimal markups, we also account for the changes in frictions implied by the counterfactual ownership, using estimates from Tables 10 and 11.

5.1 Exact hat algebra for mergers

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 the changes in brand market shares entailed by rearranging ownership and therefore altering first-order conditions for pricing. So far as we know, this is the first application of exact hat algebra (EHA) to merger analysis. Given the very low information requirements (just market shares, prices and σ must be known) this approach seems attractive as compared to methods that involve solving the full model and thus require data on the levels of A_{bn} and c_{bn} which are generally unknown. With EHA, only changes in A_{bn} and c_{bn} need to be specified and they can be obtained from the appeal and cost regressions of the previous section.

The first (and last) step in the ownership change counterfactual is to aggregate up the new brand market shares predicted by EHA to the level of firms. Initially we set $s'_{bn} = s_{bn}$, implying $\hat{s}_{bn} = 1$ and sum up the shares of the brands in the new ownership sets, $\mathcal{F}'_f$, to yield

$$\hat{S}_{fn} = \frac{\sum_{b \in \mathcal{F}'_f} \mathbb{I}_{bn} \hat{s}_{bn} s_{bn}}{S_{fn}} \tag{24}$$

The second step uses data on initial firm market shares S_{fn} and an estimate of σ to calculate the change in markups (applying a conduct assumption from equation (10). The proportional change in the Lerner index under the two alternative conduct assumptions are

$$\underbrace{\hat{L}_{fn} = \frac{\sigma - (\sigma - \eta) S_{fn}}{\sigma - (\sigma - \eta)\hat{S}_{fn} S_{fn}}}_{\text{Bertrand}}, \quad \text{and} \quad \underbrace{\hat{L}_{fn} = \frac{1 + (\sigma/\eta - 1)\hat{S}_{fn} S_{fn}}{1 + (\sigma/\eta - 1) S_{fn}}}_{\text{Cournot}} \tag{25}$$

The adjustment of the firm-level price-cost markups are

$$\hat{\mu}_{fn} = \frac{1 - L_{fn}}{1 - \hat{L}_{fn} L_{fn}} \tag{26}$$

With these markup adjustments calculated, the final step is to determine 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 $\hat{\varphi}_{bn}$. For example, consumers might like a brand better when it is under domestic ownership. The proportional change in brand-level market share is given by

$$\hat{s}_{bn} = \left(\frac{\hat{\mu}_{fn} \hat{c}_{bn}}{\hat{A}_{bn} \hat{P}_{gn}} \right)^{1-\sigma} = \left(\frac{\hat{\mu}_{fn}}{\hat{\varphi}_{bn} \hat{P}_{gn}} \right)^{1-\sigma} \tag{27}$$

with the change in country's price index given by:

$$\hat{P}_{gn} = \left(\sum_k \mathbb{I}_{kn} s_{kn} (\hat{\mu}_{kn}/\hat{\varphi}_{kn})^{1-\sigma} \right)^{\frac{1}{1-\sigma}}. \tag{28}$$

The market share adjustments $\hat{s}_{bn}$ are re-aggregated to yield changes in firm-level market shares in equation (24). The algorithm then iterates until the vector of brand-level market share changes stabilize. 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.

The presence of a set of "other" brands poses a challenge for determining counterfactual price indexes.[27] Since all the brands GMID places in the other category should have individual markets shares less than 0.1%, we model them as monopolistically competitive. Their markups are therefore fixed under CES demand, implying $\hat{\mu}_{bn} = 1$. A convenient feature of EHA counterfactual is that the aggregate market share of those brands—which we do observe and denote as s_{on}—is all we need to compute the counterfactual price index:

$$\hat{P}_{gn} = \left(s_{on} + \sum_{k \in \text{listed}_n} \mathbb{I}_{kn} s_{kn} (\hat{\mu}_{kn}/\hat{\varphi}_{kn})^{1-\sigma} \right)^{\frac{1}{1-\sigma}}. \tag{29}$$

For the listed brands, markups adjust type ownership changes. The listed brand market shares adjust according to equation (27) with the price index change given by (29). For the other brands, we assume there are no cost and appeal changes. Since markups are fixed as well, other brand market share evolves according to $\hat{s}_{on} = \hat{P}_{gn}^{\sigma-1}$.

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

$$\hat{X}_{gn} = \hat{P}_{gn}^{1-\eta}. \tag{30}$$

5.2 Outcome measures

A complete welfare calculation would be an extremely complex undertaking. While our counterfactuals provide us with a complete set of profit changes for all firm-markets, they

[27] This is related to the problem faced by Redding and Weinstein (2018), who lack detailed information on the non-traded varieties when constructing their CES price index.

do not allow us to map those outcomes to individual shareholders or taxpayers. This is because of the intricate capital structures of the many multinational firms we work with, as well as the well-known issues of how and where those profits are taxed.

To focus on what our exercise can precisely predict, we will report the following outcomes:

1. The percentage change in the price index in each product category-market $\hat{P}_{gn}$, described in equation (29).

2. The change in the Lerner index of all firms selling in n, ΔL_n.[28] Aggregate markups have been computed in several ways (see our appendix E for a comparison and how different versions of L_n can be related to the Herfindhal index in our model). In the main text, we follow Hall (2018) and use a sales-weighted Lerner index:

$$L_n = \sum_f S_{fn} L_{fn}. \tag{31}$$

3. We will also aggregate the changes in markups over markets to construct the counterfactual Lerner index at the firm level, ΔL_f. The consolidated markup L_f, shown in equation (12) as the ratio of world profits to sales, can also be expressed as

$$L_f = \sum_n \omega_{fn} L_{fn}, \tag{32}$$

where the weights are $\omega_{fn} = S_{fn}X_{gn}/S_{fw}X_{gw}$ and the w subscript denotes world. Figure 4 compares levels of L_f to those obtained from consolidated corporate accounts. In our counterfactuals we compare predicted rises in market power to those observed in the accounting data.

5.3 Restoring 2007 owners: counterfactual results

Figure 6 illustrates the model-based quantification of the impact of mergers and acquisitions occurring over the decade after 2007. The graphs in the left column contrast Herfindahl indexes with 2018 actual owners (y-axis) with simulated 2007 ownership (x-axis). The graphs on the right display changes in the price index attributed to the 2007–2018 ownership changes against changes in concentration: $\Delta H = H_{2018} - H'_{2007}$. The upper two panels show results for beer and the lower two panels show the spirits results.

[28] As noted by Syverson (2019b), rises in the aggregate Lerner index has an ambiguous effect on welfare. However, the price index impact on consumer welfare is always negative.

Figure 6: Counterfactual results: restoring the 2007 owner in 2018

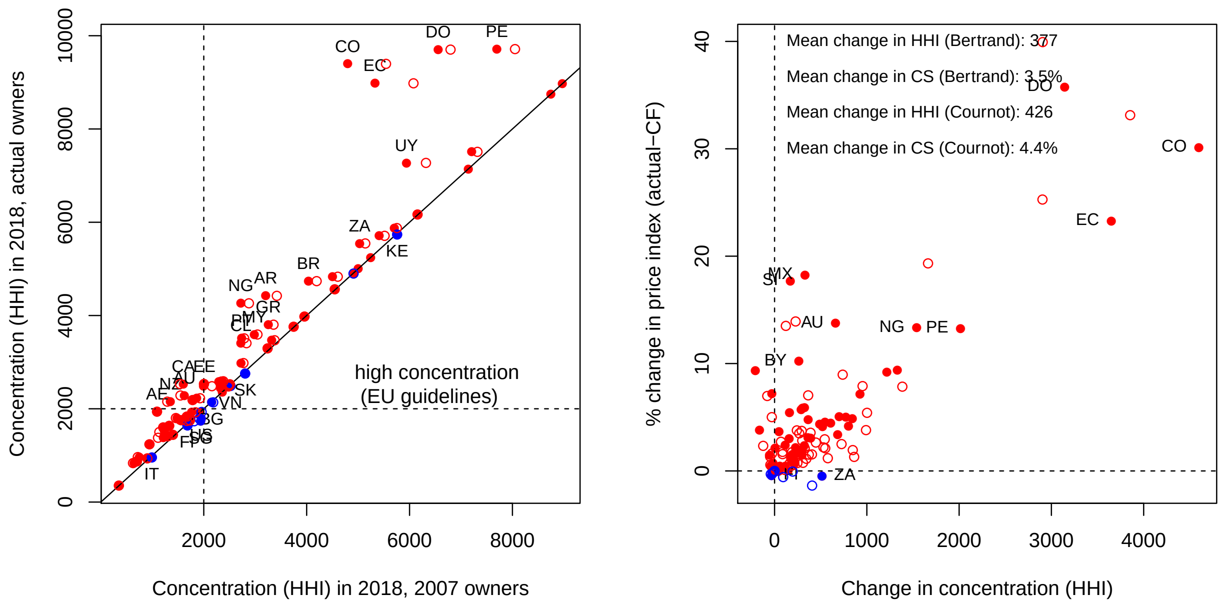

(a) Beer brand ownership changes

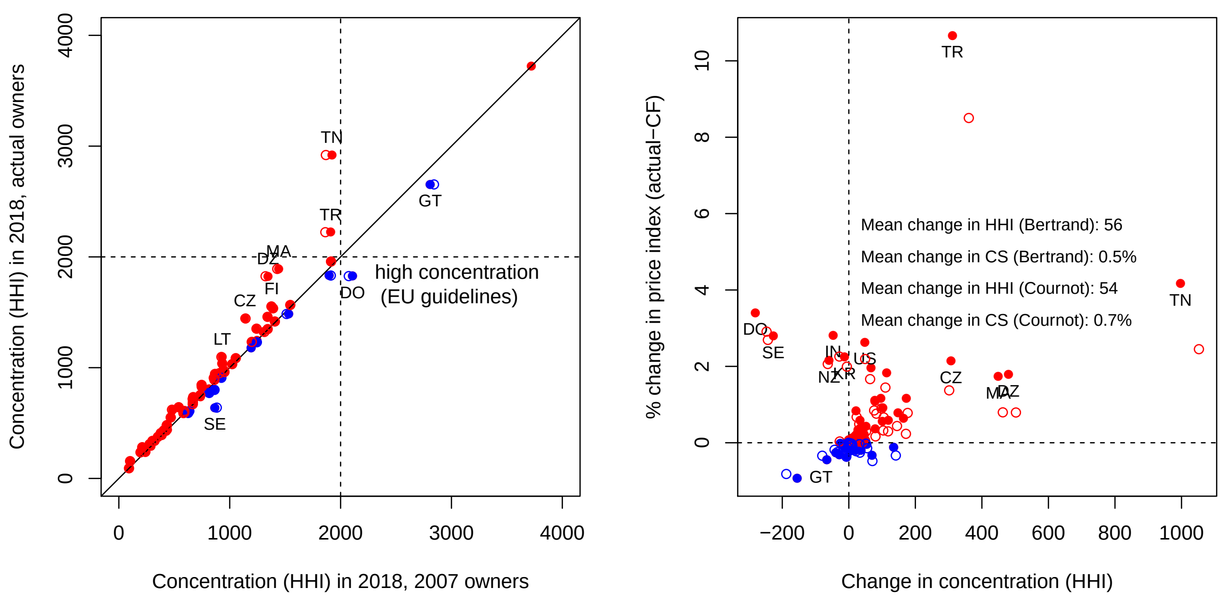

(b) Spirits brand ownership changes

Restoring 2007 owners leads to only small changes concentration for the majority of the countries. Some countries, like Canada and Australia, have concentration increases large enough to move them of the EU threshold for highly concentrated markets (shown as a horizontal dashed line). Three countries—Colombia, Peru, and the Dominican Republic—that were already highly concentrated, move to near monopoly in 2018.

The changes in the price index are very high for beer in the countries with large concentration increases. On average price indexes rise by 3 to 4%, but the countries going to nearly monopoly have an order of magnitude larger impacts. On average Cournot conduct leads to mergers that have higher concentration and harm to consumers but at very high levels of concentration, the Bertrand assumption generates greater losses. The intuition for this can be found in the convexity of the Lerner index as a function of market share under Bertrand conduct, as illustrated in Figure 4(a).

The counterfactuals underlying Figure 6 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 changes in φ^F_{bn} are taken from the HQ frictions coefficients in columns (3) for Bertrand and (5) for Cournot of Tables 10 and 11 in the Appendix.

To isolate the pure market power effects, Appendix D displays the case where ownership changes do not alter brand type ($\hat{\varphi}^F = 1$). That case is interesting because it confirms the approximation result in Proposition 5 of Nocke and Schutz (2018a). Namely, among the 10 or so countries that experience modest rises in concentration, the rise in the price index is roughly linear in the change in concentration (ΔH_n).

The counterfactual results for M&A between spirits firms, shown in the lower two panels Figure 6, single out Turkey as the only market with major increases in concentration and the price index. The trigger for this rise in market power is Diageo's acquisition of the owner of Yeni Raki, the most popular spirit in the country. In other markets spirits mergers had more modest effects, leading to an average price increase across the 75 markets of 0.6% (Bertrand) to 0.7% (Cournot).

Figure 7 graphs on the vertical axis the changes in the Lerner index based on the changes in firm market shares that were actually observed from 2007 to 2018. On the x-axis, the figure displays the changes in market power that the counterfactual simulation attributes to all the ownership changes over the last decade. Only 7 countries out of 75 were unaffected by acquisitions or divestments during the period. They appear as red circles on the y-axis. Of the 68 countries with ownership changes, 61 were market-power increasing. Among the cases where divestiture dominated acquisitions, South Africa saw the biggest decline in market power. This came as spillover from EU competition policy:

Figure 7: Effects of ownership changes 2007–18 on country-level markups

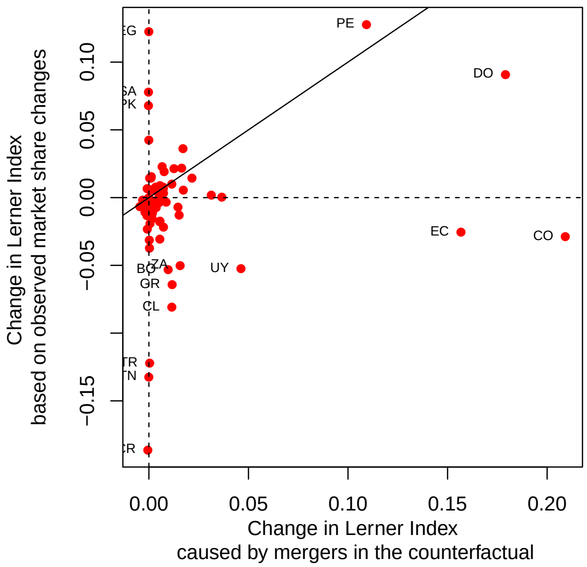

(a) Beer: All countries

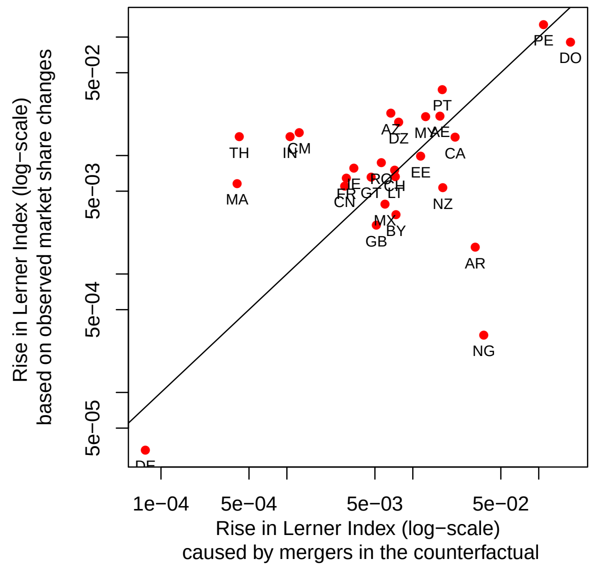

(b) Countries with increasing market power

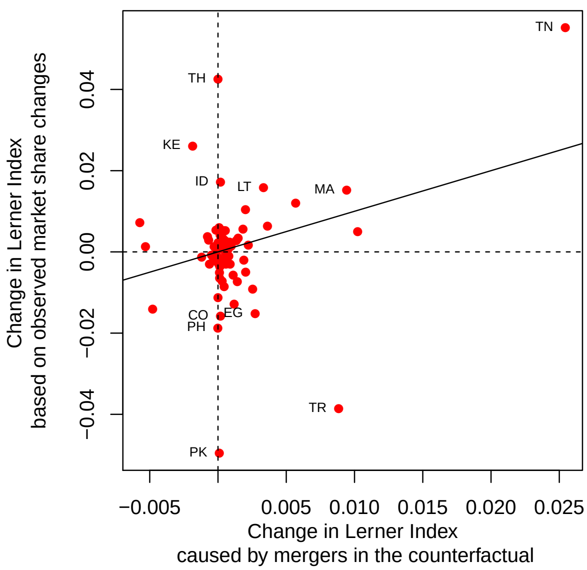

(a) Spirits: All countries

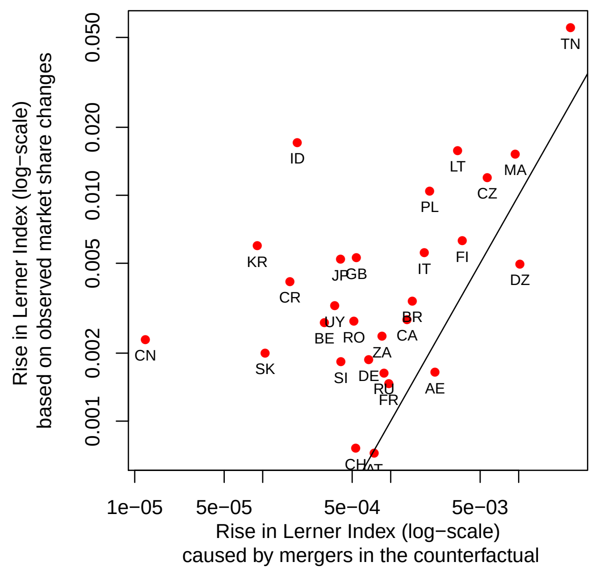

(b) Countries with increasing market power

AB InBev was compelled to divest Pilsner Urquell and Miller when it acquired SABMiller.

The EHA method hold firm types constant at their 2018 values in the x-axis changes, but implicitly includes the changes in firm type between 2007 and 2018 on the y-axis, along with rises in firm market shares generated by mergers. Thus, the y-axis combines merger effects with superstar/convergence effects. Superstar effects are increases in concentration due to rising market shares of the already large firms. Convergence effects occur when smaller firms expand market shares at the expense of the larger firms.[29] The y-axis ΔL_n also reflect net entry between 2007 and 2018 (a convergence effect operating on the extensive margin).

The 45-degree line in Figure 7 divides the countries in our data set into those where superstar effects boost market power (raise L_n) to a larger extent than predicted by M&A alone (above the line) and those where convergence effects dominate M&A and suppress market power (below the line). Along the line, we see Argentina (AR) and Peru (PE) as the two countries where the movement in Lerner index based on observed changes in market share corresponds to the ΔL_n predicted by the mergers.

Under the 45-degree, we see that 60% of countries (45 out of 70) had firm market share movements that constrict market power (convergence effects). Especially notable are the five Latin American countries where mergers were predicted to lead to 4-point or greater increases in L_n but the actual changes in market shares were much smaller. Uruguay is a case where mergers were predicting a 5-point increase in L_n but in fact L_n fell by about 7 points. This is because AB Inbev's brands fell in market share from 90% to 83% despite adding global giant Corona. Underlying the poor performance were declines in the local stars: Pilsen, the market leader went from 53% to 33%. Meanwhile, weak brands strengthened ("Others" double their collective market share) and a divestment agreement gave Miller to MolsonCoors, creating a new entrant in the Uruguay beer market.

There are 15 countries where superstar effects add to the market power increase attributed to mergers. The most striking case is Egypt where no mergers took place but market share increases lead to a 11-point increase in market power (from an already high index of 33%). Heineken NV in Egypt fits the superstar firm story perfectly. Its two main brands increase: Stella (a local star) rises from 39% to 42% and global giant Heineken (brand) went from 8.8 to 14.6%. Heineken, as a corporation, increased its share from 63% to 81%. Meanwhile, the share of others crashed from 25% to 3.5%.

Figure 7 (a) seems to have a bad fit. But this is not the right way to interpret it. The dominant feature of the graph is that even among countries where there was substantial

[29] We borrow the superstar term from Autor et al. (2017) and convergence from the cross-country growth literature.

M&A (predicting rising market power), most countries had declines in market shares of the dominant firms leading to lower marker power (L_n). That means convergence, not superstar effects, prevailed in the aggregate. Figure 7 (b) zooms in on the truncated set of countries with mergers and with increases in L_n (based on market share changes). The correlation (in levels) between the change in L_n attributable to mergers and the L_n increase caused by market share changes from all sources is 0.86. The combined message of the two plots is that where market power was rising, M&A was the driving force behind the increase (except in two countries). Furthermore, had competition policy been more active, market power would have fallen even more from 2007 to 2018 than the decreases shown in panel (a) for many South American countries.

Another attraction of using the share-weighted Lerner index is attractive as a measure of market power is that it can also be calculated at the firm level as well as the market level. The firm-level L_f is an index of markups weighted by the share of the firm's sales that comes from each market. Thus, L_f is high when the firm has high market share in the markets that contribute importantly to its global revenues.

Figure 8: Effects of ownership changes 2007–16 on firm-level markups

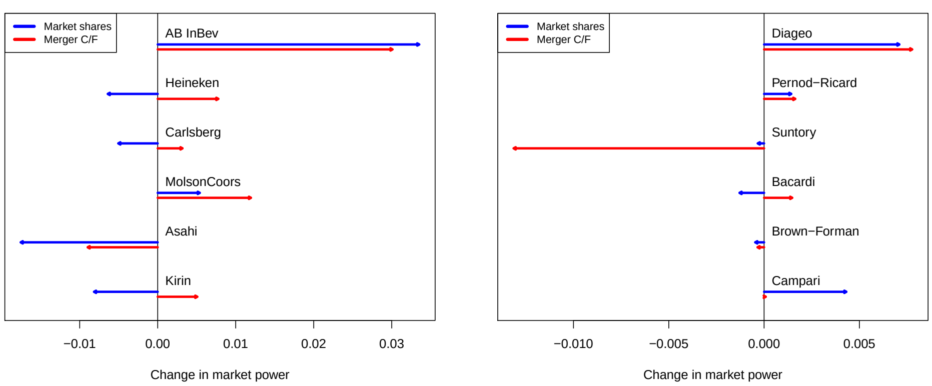

(a) Beer (b) Spirits

Figure 8 shows with the red arrows the changes in firm-level market power brought about by M&A and divestments from 2007–2018 for the six largest beer makers at the end of 2018. We show (in blue) the changes in L_f generated by actual changes in market share from 2007 to 2018. As with the market-level measure of market power, this calculation combines extensive margin changes in firm type (coming from expansion of the brand portfolios, $\mathcal{F}_f$) and changes in the appeal or costs (summarized in φ_{bn}) of the *incumbent* brands in each market where the brands are sold. Thus a firm could see its type decline—even as it added brands—if the φ_{bn} in 2018 for important brand-markets were lower than 2007 levels. This is precisely what happens to Heineken and Carlsberg. Despite numerous acquisitions in multiple markets (as shown in Figure 2), the losses of market share in strongholds of those two firms (notably Spain, Poland and Greece for Heineken, and all Nordic countries for Carlsberg) dominated the gains in markets entered through acquisitions.

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 complete Japan focus, where their market shares were dominant (40% and 31% market shares, respectively), to diversified positions where lower market share brands contribute substantially to total sales.[30] 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.

AB InBev's performance in Figure 8 is a mirror-image for Asahi: Acquired brands explain almost all of the change in market power but now L_f is rising, not falling. AB InBev's market power index, already highest of the top six beer makers in 2007, rises a further four percentage points over the next decade.

5.4 Undoing forced divestitures: counterfactual results

Non-academic narratives frequently portray competition authorities as passively permitting monopolization. On the other hand, Gutierrez and Philippon (2018) distinguish the EU case as being strongly affected by regulation unlike the more lax US policy environment. We observe that in the beer industry competition authorities on both sides of the Atlantic have forced divestitures to avoid concentration and even multi-market coordination effects.[31]

AB InBev has been 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. Second, when it bought the Modelo Group, it had to divest the US-market rights of Corona several other brands to Constellation Brands (a company mainly active in wine). The acquisition of SAB Miller in 2016 triggered forced divestitures in the US, EU, and China. Specifically, a package of popular EU brands was sold to Asahi, all the Miller brands were sold to MolsonCoors, and AB InBev's minority share of China Resources was sold to its Chinese partner.

Our model and data are well-suited to evaluating the efficacy of these divestiture by

[30] Both firms obtain 98% of sales from Japan it 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.

[31] See European Commission decision.

Table 7: What if antitrust authorities had been more permissive?

Country	$\hat{P}_{gn}$ with $\hat{\varphi}^F_{f(b,t)} = 1$		$\hat{P}_{gn}$ with $\hat{\varphi}^F_{f(b,t)} \neq 1$	
	Bertrand	Cournot	Bertrand	Cournot
United States	4.23	5.88	5.65	7.20
United Arab Emirates	1.13	1.91	1.09	1.87
Netherlands	1.04	2.04	0.72	1.61
Hungary	1.03	1.83	-0.01	0.67
Italy	0.79	1.58	0.25	1.02
Czechia	0.54	0.78	-1.10	-1.01
Slovakia	0.20	0.34	-1.11	-1.11
Poland	0.00	0.00	-1.26	-1.40

Note: 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.

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 a percent are displayed in Table 7. Sorted in descending order by the price change for Bertrand without cost adjustments , the table also includes prices changes for Cournot, and both Bertrand and Cournot imposing the adjustment to $\varphi^F_{f(b,t)}$ predicted in our regression analysis for beer.

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 seven percent higher. The highest price increase occurs under Cournot competition when the frictions are higher because the HQ of Miller brands switches from the US to Belgium. The EU commission's intervention protected consumers from non-negligible increases in market power in Hungary, Netherlands, and Italy. In the first case, 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 (38 markets), Leffe (10 markets) and Becks (34 markets), which collectively held 7% of the market. In Italy AB InBev brands (led by Becks at 6%) accounted for 13% of market in 2016, similar to Asahi's 14% (8% of which was Peroni).

The case of United Arab Emirates 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. The

UAE 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 market situations in Czechia, Slovakia, and Poland exemplify the unintended consequences of divestiture to a remote owner. In those countries, the simulation predicts minimal (or zero in the case of Poland) prices rises due to market power.[32] However, the move of HQ from Brussels to Tokyo increases frictions by enough to raise price index of beer by 1.3 to 1.4%. The potential costs of distance between market and headquarters is an issue that can only be quantified by combining data from multiple markets as we do in this paper.

5.5 Forcing counterfactual divestitures

The US and EU competition authorities' policy of forced divestitures as pre-conditions for approving InBev and its successor's acquisitions appears to have resulted in sizeable savings in the US and modest savings in several other countries as well. This raises the question of whether competition agencies in other countries could have achieved similar consumer savings by emulating the US/EU approach. The counterfactual reported in Table 8 reassigns the global rights to Labatt brands to FIFCO (the company that later bought North American Breweries), 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 is tantamount to placing the pricing decisions for these brands under independent control.

The largest gains would accrue to consumers in three Andean countries where SABMiller had acquired the main local star brands. Forcing divestiture 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.

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 8 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

[32] In Poland, AB InBev retained no other brands (above the GMID threshold) after the divestiture. This implies no change in markups due to ownership changes considered in isolation. The EU Commission justified the divestiture of the Polish brands based on concerns based on multi-market contacts.

Table 8: What if antitrust authorities had followed EU/US lead?

Country	$\hat{P}_{gn}$ with $\hat{\varphi}^F_{f(b,t)} = 1$		$\hat{P}_{gn}$ with $\hat{\varphi}^F_{f(b,t)} \neq 1$	
	Bertrand	Cournot	Bertrand	Cournot
Colombia	-30.21	-25.87	-30.08	-25.71
Ecuador	-25.26	-22.69	-25.14	-22.54
Peru	-19.59	-14.05	-19.45	-13.76
Uruguay	-10.12	-11.54	-10.22	-11.66
Dominican Republic	-7.05	-4.18	-7.33	-4.40
Canada	-2.65	-5.50	-3.45	-6.43
Argentina	-2.24	-4.22	-2.21	-4.16
Australia	-1.97	-4.32	-1.89	-4.20
United Arab Emirates	-1.72	-3.77	-1.50	-3.48
Bolivia	-1.63	-2.12	-1.67	-2.18
Mexico	-1.35	-2.94	-2.47	-4.21
Chile	-1.16	-2.71	-1.22	-2.79
South Africa	-1.11	-2.05	-0.67	-1.47
Guatemala	-0.66	-1.50	-0.84	-1.75
India	-0.37	-0.95	-0.42	-1.02

Note: The table reports the effect of forcing divestitures on the percent change in the price index for beer in each country in 2018.

Mexican market.

The price reductions reported in Table 8 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 SAB Miller brands it retained.[33] Our counterfactuals suggest the main benefits were near monopolization of several Latin American beer markets.

6 Conclusion

In the beer and spirits industries, a small group of firms, headquartered in a handful of countries, have been acquiring brands all over the world. 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

[33]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).

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 performance—extracted from data on market shares—is, for the most part, invariant to the identity of the owner. That is, firm fixed effects explain just 3% of cost-adjusted appeal in the beer industry and 4% in spirits. The reduction in R^2 from excluding firm fixed effects is close to zero. There is one way that ownership *does* affect brand performance, however. In the spirits industry, and to a lesser extent, the beer industry, we estimate that brands operate at lower costs in the countries where their owners are headquartered. These results suggest there is cost penalty from foreign acquisitions with little in the way of synergies. 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.

Rises in concentration at the world level can substantially overstate the changes in concentration in specific countries. To see this, consider a firm that acquires local monopolists in two countries. The MNCs world market share rises but the market structure within those markets is unchanged. The crucial condition that makes acquisitions profitable for the MNC—and harmful for the local consumers—is whether mergers combine global giants and local stars in the same market.

Cross-country comparisons in our counterfactuals quantify the beneficial role of competition policy towards mergers. The divestitures forced by the US and EU led to significant consumer savings, especially in the US. Canada and Australia could have achieved similar savings by imposing divestitures along the same lines. The greatest potential for the use of stricter competition policy would be in Latin America, where counterfactuals reveal that consumer prices increases of over 20% could have been avoided in some countries.

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 (2017). The elusive pro-competitive effects of trade. *The Review of Economic Studies*.

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 (2017, May). The fall of the labor share and the rise of superstar firms. Working Paper 23396, National Bureau of Economic Research.

Basu, S. (2019, August). Are price-cost markups rising in the United States? a discussion of the evidence. *Journal of Economic Perspectives 33*(3), 3–22.

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, August). 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, October). Evidence for the effects of mergers on market power and efficiency. Working Paper 22750, National Bureau of Economic Research.

Brander, J. A. and P. R. Krugman (1983). A 'reciprocal dumping' model of international trade. *Journal of International Economics 15*(3), 313–321.

Coşar, K., P. Grieco, S. Li, and F. Tintelnot (2018). What drives home market advantage? *Journal of International Economics 110*, 135–150.

Crozet, M., K. Head, and T. Mayer (2012). Quality sorting and trade: Firm-level evidence for french wine. *The Review of Economic Studies 79*(2), 609–644.

De Loecker, J. and J. Eeckhout (2017). The rise of market power and the macroeconomic implications. Working Paper 23687, National Bureau of Economic Research.

De Loecker, J. and J. Eeckhout (2018, June). Global market power. Working Paper 24768, National Bureau of Economic Research.

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, December). 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.

Grullon, G., Y. Larkin, and R. Michaely (2019, 04). Are US Industries Becoming More Concentrated?*. *Review of Finance 23*(4), 697–743.

Gutierrez, G. and T. Philippon (2018, June). How EU markets became more competitive than US markets: A study of institutional drift. Working Paper 24700, National Bureau of Economic Research.

Hall, R. E. (2018, November). Using empirical marginal cost to measure market power in the US economy. Working Paper 25251, National Bureau of Economic Research.

Head, K. and T. Mayer (2018). Poor substitutes? Counterfactual methods in IO and trade compared. mimeo.

Head, K. and T. Mayer (2019, September). Brands in motion: How frictions shape multinational production. *American Economic Review 109*(9), 3073–3124.

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.

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.

Kugler, M. and E. Verhoogen (2011, 11). Prices, plant size, and product quality. *The Review of Economic Studies 79*(1), 307–339.

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.

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.

Redding, S. J. and D. E. Weinstein (2018). Accounting for trade patterns. Discussion Paper 12446, Center for Economic Policy Research.

Syverson, C. (2019a, August). 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*.

Zhelobodko, E., S. Kokovin, M. Parenti, and J.-F. Thisse (2012). Monopolistic competition: Beyond the constant elasticity of substitution. *Econometrica 80*(6), 2765–2784.

A Entry and exit rates for brands and markets

In this section, we document entry rates for each of the beverage categories. We investigate whether another extensive margin is also common: the entry margin, through which firms add or drop brands in selected markets or altogether. This is done in Table 9.

Table 9: Adding and dropping brands in markets and overall: Beer and Spirits

Sample frame	Add rate (in percent)	Drop rate (in percent)
Beer		
Brand-level births and deaths:		
All brand/years (18,063 obs.)	3.07	2.34
Brands changing owners: before	NA	2.21
Brands changing owners: after	NA	2.49
Brands added/dropped in a market:		
All brand/market/years (1,498,802 obs.)	0.06	2.63
Continuing brands	0.02	0.85
Brands changing owners: before	0.02	0.32
Brands changing owners: after	0.03	1.90
Spirits		
(a) Brand-level births and deaths:		
All brand/years (25,601 obs.)	1.62	1.69
Brands changing owners: before	NA	1.69
Brands changing owners: after	NA	1.50
(b) Brands added/dropped in a market:		
All brand/market/years (1,919,019 obs.)	0.04	1.56
Continuing brands	0.02	0.61
Brands changing owners: before	0.01	0.51
Brands changing owners: after	0.04	1.93

B Additional regression results

B.1 Beer and Spirits estimated separately

Table 10: Explaining appeal, cost, and markups for Beer

	$\ln s_{bn}$	$\ln A_{bn}$	$\ln \varphi_{bn}$	$\ln \mu_{bn}$	$\ln \varphi_{bn}$	$\ln \mu_{bn}$
			Bertrand		Cournot	
home	1.480^{a}	0.211^{a}	0.452^{a}	0.028^{a}	0.479^{a}	0.056^{a}
	(0.173)	(0.065)	(0.053)	(0.007)	(0.056)	(0.012)
distance	-0.250^{a}	-0.034	-0.079^{a}	-0.007^{b}	-0.084^{a}	-0.013^{a}
	(0.060)	(0.023)	(0.018)	(0.003)	(0.019)	(0.004)
common language	0.351^{a}	0.030	0.104^{a}	0.003	0.106^{b}	0.005
	(0.135)	(0.056)	(0.040)	(0.005)	(0.042)	(0.009)
home (HQ)	0.182	0.051	0.099^{c}	0.047^{a}	0.131^{b}	0.079^{a}
	(0.197)	(0.069)	(0.059)	(0.009)	(0.063)	(0.015)
distance (HQ)	-0.075	-0.019	-0.019	0.003	-0.022	-0.001
	(0.054)	(0.022)	(0.016)	(0.003)	(0.017)	(0.004)
com. lang. (HQ)	-0.149	-0.014	-0.047	-0.005	-0.046	-0.003
	(0.112)	(0.055)	(0.034)	(0.007)	(0.036)	(0.010)
Observations	34,653	34,653	34,653	34,653	34,653	34,653
R^2	0.758	0.792	0.737	0.887	0.745	0.867

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

Table 11: Explaining appeal, cost, and markups for Spirits

	$\ln s_{bn}$	$\ln A_{bn}$	$\ln \varphi_{bn}$	$\ln \mu_{bn}$	$\ln \varphi_{bn}$	$\ln \mu_{bn}$
			Bertrand		Cournot	
home	0.652[a]	0.214[b]	0.277[a]	0.003	0.279[a]	0.005
	(0.157)	(0.105)	(0.067)	(0.002)	(0.068)	(0.004)
distance	−0.075[c]	0.065[b]	−0.032[c]	−0.0004	−0.032[c]	−0.001
	(0.045)	(0.030)	(0.019)	(0.001)	(0.019)	(0.001)
common language	−0.042	−0.078	−0.019	−0.001	−0.019	−0.002
	(0.093)	(0.066)	(0.039)	(0.001)	(0.040)	(0.002)
home (HQ)	0.440[a]	0.116	0.211[a]	0.026[a]	0.233[a]	0.048[a]
	(0.131)	(0.084)	(0.057)	(0.003)	(0.058)	(0.005)
distance (HQ)	0.067[c]	0.023	0.029[c]	0.001	0.030[c]	0.002
	(0.039)	(0.026)	(0.017)	(0.001)	(0.017)	(0.001)
com. lang. (HQ)	0.165[b]	0.063	0.075[b]	0.006[b]	0.079[b]	0.010[b]
	(0.070)	(0.047)	(0.030)	(0.003)	(0.031)	(0.005)
Observations	60,592	60,592	60,592	60,592	60,592	60,592
R^2	0.578	0.619	0.550	0.864	0.553	0.846

Table 12: Explaining appeal, type, and markups for Beer, without firm fixed effects

	$\ln s_{bn}$	$\ln A_{bn}$	$\ln \varphi_{bn}$	$\ln \mu_{bn}$	$\ln \varphi_{bn}$	$\ln \mu_{bn}$
			Bertrand		Cournot	
home	1.511[a]	0.214[a]	0.475[a]	0.043[a]	0.512[a]	0.080[a]
	(0.162)	(0.058)	(0.050)	(0.007)	(0.052)	(0.012)
distance	−0.216[a]	−0.037[c]	−0.067[a]	−0.005[b]	−0.072[a]	−0.010[a]
	(0.056)	(0.021)	(0.017)	(0.003)	(0.018)	(0.004)
common language	0.375[a]	0.041	0.115[a]	0.008	0.122[a]	0.015[c]
	(0.129)	(0.054)	(0.038)	(0.005)	(0.040)	(0.009)
home (HQ)	0.061	0.034	0.033	0.015[b]	0.042	0.024[b]
	(0.138)	(0.048)	(0.042)	(0.006)	(0.044)	(0.011)
distance (HQ)	−0.087[b]	−0.015	−0.026[b]	−0.001	−0.029[a]	−0.004[c]
	(0.035)	(0.014)	(0.010)	(0.002)	(0.011)	(0.003)
com. lang. (HQ)	−0.158[c]	−0.013	−0.059[b]	−0.014[b]	−0.067[b]	−0.021[b]
	(0.095)	(0.045)	(0.029)	(0.006)	(0.030)	(0.010)
Observations	34,653	34,653	34,653	34,653	34,653	34,653
R^2	0.750	0.785	0.729	0.876	0.736	0.853

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

Table 13: Explaining appeal, cost, and markups for Beer, largest connected set

	$\ln s_{bn}$	$\ln A_{bn}$	$\ln \varphi_{bn}$	$\ln \mu_{bn}$	$\ln \varphi_{bn}$	$\ln \mu_{bn}$
			Bertrand		Cournot	
home	1.432^{a}	0.208^{a}	0.442^{a}	0.032^{a}	0.471^{a}	0.061^{a}
	(0.178)	(0.067)	(0.054)	(0.008)	(0.058)	(0.012)
distance	−0.270^{a}	−0.041^{c}	−0.083^{a}	−0.006^{b}	−0.089^{a}	−0.012^{b}
	(0.061)	(0.023)	(0.019)	(0.003)	(0.020)	(0.005)
common language	0.269^{c}	0.007	0.079^{c}	0.002	0.080^{c}	0.003
	(0.139)	(0.058)	(0.041)	(0.005)	(0.043)	(0.008)
home (HQ)	0.139	0.087	0.106	0.067^{a}	0.150^{b}	0.110^{a}
	(0.220)	(0.078)	(0.068)	(0.012)	(0.074)	(0.021)
distance (HQ)	−0.006	0.006	0.004	0.006^{b}	0.002	0.004
	(0.055)	(0.023)	(0.016)	(0.003)	(0.017)	(0.004)
com. lang. (HQ)	−0.132	0.009	−0.041	−0.003	−0.037	0.001
	(0.122)	(0.058)	(0.037)	(0.007)	(0.039)	(0.012)
Observations	24,680	24,680	24,680	24,680	24,680	24,680
R^2	0.726	0.738	0.703	0.874	0.707	0.840

Table 14: Explaining appeal, type, and markups for Spirits, without firm fixed effects

	$\ln s_{bn}$	$\ln A_{bn}$	$\ln \varphi_{bn}$	$\ln \mu_{bn}$	$\ln \varphi_{bn}$	$\ln \mu_{bn}$
			Bertrand		Cournot	
home	0.611^{a}	0.196^{c}	0.262^{a}	0.005^{b}	0.266^{a}	0.009^{b}
	(0.151)	(0.101)	(0.065)	(0.002)	(0.066)	(0.004)
distance	−0.074^{c}	0.059^{b}	−0.031^{c}	−0.0004	−0.032^{c}	−0.001
	(0.043)	(0.028)	(0.018)	(0.001)	(0.019)	(0.001)
common language	−0.044	−0.082	−0.020	−0.002	−0.021	−0.002
	(0.090)	(0.064)	(0.038)	(0.001)	(0.039)	(0.002)
home (HQ)	0.424^{a}	0.101	0.198^{a}	0.020^{a}	0.215^{a}	0.037^{a}
	(0.118)	(0.075)	(0.051)	(0.003)	(0.053)	(0.005)
distance (HQ)	0.064^{c}	0.023	0.028^{c}	0.001	0.028^{c}	0.001
	(0.036)	(0.023)	(0.015)	(0.001)	(0.015)	(0.001)
com. lang. (HQ)	0.148^{b}	0.059	0.067^{b}	0.005^{c}	0.071^{b}	0.009^{b}
	(0.066)	(0.044)	(0.029)	(0.003)	(0.030)	(0.004)
Observations	60,592	60,592	60,592	60,592	60,592	60,592
R^2	0.571	0.614	0.542	0.849	0.545	0.829

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

Table 15: Explaining appeal, cost, and markups for Spirits, largest connected set

	$\ln s_{bn}$	$\ln A_{bn}$	$\ln \varphi_{bn}$	$\ln \mu_{bn}$	$\ln \varphi_{bn}$	$\ln \mu_{bn}$
			Bertrand		Cournot	
home	0.648^a	0.221^c	0.276^a	0.003	0.278^a	0.006
	(0.173)	(0.117)	(0.074)	(0.002)	(0.074)	(0.004)
distance	−0.009	0.088^a	−0.003	0.0003	−0.003	0.001
	(0.044)	(0.031)	(0.019)	(0.001)	(0.019)	(0.001)
common language	−0.015	−0.068	−0.007	−0.001	−0.007	−0.001
	(0.094)	(0.067)	(0.040)	(0.001)	(0.041)	(0.002)
home (HQ)	0.331^b	0.081	0.171^a	0.031^a	0.196^a	0.057^a
	(0.145)	(0.091)	(0.062)	(0.003)	(0.064)	(0.006)
distance (HQ)	0.062	0.015	0.027	0.001	0.028	0.001
	(0.041)	(0.027)	(0.017)	(0.001)	(0.018)	(0.001)
com. lang. (HQ)	0.164^b	0.067	0.075^b	0.006^b	0.079^b	0.010^b
	(0.074)	(0.050)	(0.032)	(0.003)	(0.033)	(0.005)
Observations	40,254	40,254	40,254	40,254	40,254	40,254
R^2	0.472	0.525	0.451	0.792	0.457	0.798

B.2 Results pooling Beer and Spirits

Table 16: Pooled 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	1.030^{a}	0.212^{a}	0.359^{a}	0.023^{a}	0.379^{a}	0.043^{a}
	(0.125)	(0.068)	(0.047)	(0.004)	(0.049)	(0.006)
distance	-0.111^{a}	0.029	-0.040^{a}	-0.002^{c}	-0.042^{a}	-0.003^{b}
	(0.035)	(0.020)	(0.014)	(0.001)	(0.014)	(0.002)
common language	0.059	−0.054	0.013	0.0004	0.014	0.001
	(0.076)	(0.049)	(0.030)	(0.002)	(0.031)	(0.003)
home (HQ)	0.268^{a}	0.080	0.135^{a}	0.018^{a}	0.148^{a}	0.030^{a}
	(0.093)	(0.051)	(0.037)	(0.003)	(0.038)	(0.006)
distance (HQ)	0.009	0.010	0.009	0.0002	0.008	−0.001
	(0.027)	(0.016)	(0.011)	(0.001)	(0.011)	(0.001)
com. lang. (HQ)	0.082	0.046	0.038^{c}	0.001	0.039^{c}	0.002
	(0.056)	(0.035)	(0.023)	(0.003)	(0.024)	(0.004)
Observations	95,245	95,245	95,245	95,245	95,245	95,245
R^2	0.651	0.649	0.588	0.891	0.596	0.845

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 headquarter country. Significance levels: 1% (a), 5% (b), and 10% (c).

Table 17: Pooled 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	1.063^a	0.238^a	0.366^a	0.019^a	0.382^a	0.036^a
	(0.149)	(0.079)	(0.055)	(0.004)	(0.057)	(0.007)
distance	-0.085^b	0.049^b	-0.027^c	-0.001	-0.029^c	-0.003
	(0.038)	(0.023)	(0.015)	(0.001)	(0.015)	(0.002)
common language	0.053	-0.053	0.012	0.0004	0.012	0.001
	(0.080)	(0.052)	(0.032)	(0.002)	(0.033)	(0.003)
home (HQ)	0.245^b	0.086	0.149^a	0.041^a	0.179^a	0.071^a
	(0.121)	(0.066)	(0.048)	(0.005)	(0.050)	(0.009)
distance (HQ)	0.041	0.012	0.020	0.002^b	0.020	0.002
	(0.034)	(0.021)	(0.014)	(0.001)	(0.014)	(0.001)
com. lang. (HQ)	0.108^c	0.058	0.052^c	0.004	0.056^b	0.008^c
	(0.065)	(0.041)	(0.027)	(0.003)	(0.028)	(0.005)
Observations	64,934	64,934	64,934	64,934	64,934	64,934
R^2	0.598	0.568	0.519	0.876	0.527	0.826

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 headquarter country. Significance levels: 1% (a), 5% (b), and 10% (c).

B.3 Results pooling 7 Beverages

Table 18: 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	1.003^a	0.263^a	0.394^a	0.018^a	0.408^a	0.033^a
	(0.098)	(0.056)	(0.040)	(0.003)	(0.041)	(0.005)
distance	-0.162^a	0.001	-0.055^a	-0.001	-0.056^a	-0.002
	(0.029)	(0.017)	(0.012)	(0.001)	(0.012)	(0.001)
common language	0.122^c	-0.022	0.039	0.001	0.041	0.003
	(0.064)	(0.038)	(0.026)	(0.002)	(0.026)	(0.003)
home (HQ)	0.359^a	0.136^a	0.173^a	0.021^a	0.190^a	0.038^a
	(0.081)	(0.045)	(0.033)	(0.003)	(0.034)	(0.005)
distance (HQ)	0.027	0.016	0.014	-0.0001	0.013	-0.001
	(0.026)	(0.015)	(0.010)	(0.001)	(0.011)	(0.001)
com. lang. (HQ)	0.140^a	0.063^b	0.065^a	0.003	0.068^a	0.006^c
	(0.053)	(0.032)	(0.022)	(0.002)	(0.023)	(0.003)
Observations	170,496	170,496	170,496	170,496	170,496	170,496
R^2	0.737	0.700	0.669	0.941	0.674	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 headquarter country. Significance levels: 1% (a), 5% (b), and 10% (c).

Table 19: 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	1.024^a	0.282^a	0.408^a	0.023^a	0.426^a	0.041^a
	(0.092)	(0.051)	(0.038)	(0.003)	(0.039)	(0.005)
distance	-0.149^a	-0.005	-0.052^a	-0.0001	-0.053^a	-0.001
	(0.027)	(0.015)	(0.011)	(0.001)	(0.011)	(0.001)
common language	0.129^b	-0.019	0.042^c	0.0002	0.043^c	0.001
	(0.062)	(0.036)	(0.025)	(0.002)	(0.026)	(0.003)
home (HQ)	0.287^a	0.092^b	0.128^a	0.010^a	0.136^a	0.019^a
	(0.070)	(0.037)	(0.029)	(0.002)	(0.030)	(0.004)
distance (HQ)	0.013	0.016	0.009	-0.0002	0.009	-0.001
	(0.021)	(0.012)	(0.009)	(0.001)	(0.009)	(0.001)
com. lang. (HQ)	0.101^b	0.043	0.046^b	0.003	0.048^b	0.005
	(0.049)	(0.029)	(0.020)	(0.002)	(0.021)	(0.003)
Observations	170,496	170,496	170,496	170,496	170,496	170,496
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 headquarter country. Significance levels: 1% (a), 5% (b), and 10% (c).

Table 20: Pooled regressions, 7 beverages, largest connected set

			Bertrand		Cournot	
	$\ln s_{bn}$	$\ln A_{bn}$	$\ln \varphi_{bn}$	$\ln \mu_{bn}$	$\ln \varphi_{bn}$	$\ln \mu_{bn}$
home	1.127^{a}	0.299^{a}	0.440^{a}	0.021^{a}	0.456^{a}	0.037^{a}
	(0.125)	(0.068)	(0.050)	(0.003)	(0.051)	(0.006)
distance	-0.123^{a}	0.016	-0.043^{a}	-0.00004	-0.044^{a}	-0.001
	(0.033)	(0.019)	(0.013)	(0.001)	(0.013)	(0.001)
common language	0.095	-0.032	0.032	0.001	0.034	0.002
	(0.069)	(0.043)	(0.028)	(0.002)	(0.029)	(0.003)
home (HQ)	0.213^{b}	0.083	0.119^{a}	0.036^{a}	0.146^{a}	0.063^{a}
	(0.108)	(0.058)	(0.043)	(0.005)	(0.045)	(0.008)
distance (HQ)	0.053^{c}	0.024	0.023^{c}	0.0003	0.022	-0.001
	(0.032)	(0.020)	(0.014)	(0.001)	(0.014)	(0.001)
com. lang. (HQ)	0.115^{c}	0.069^{c}	0.057^{b}	0.004^{c}	0.061^{b}	0.008^{c}
	(0.063)	(0.040)	(0.027)	(0.003)	(0.028)	(0.004)
Observations	92,048	92,048	92,048	92,048	92,048	92,048
R^2	0.675	0.589	0.586	0.899	0.593	0.863

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 headquarter country. Significance levels: 1% (a), 5% (b), and 10% (c).

B.4 Correlations of brand and firm fixed effects, with low mobility bias

Andrews et al. (2008) have shown that the low mobility bias was an important issue in employer-employee data, which was even able to change the sign of the correlation between fixed effects, therefore reversing the finding of whether there is positive or negative assortative matching. They recommend to concentrate on the set of movers (brands that change ownership in our context), and "high-mobility" firms (firms that have at least ten movements of brands over the period) to reduce that bias, which is what we do in the main text.

In this appendix, we report that same correlations, when the regressions do not apply the two above restrictions to the estimating sample. As found by labor economists, the patterns of correlations now find stronger support for *negative* assortative matching: all correlations between brands and firm fixed effects are now negative and larger in absolute value, for both beer and spirits.

Table 21: Correlations between fixed effects

	Brand			Firm		
Dep. var.:	share (s_{bn})	appeal (A_{bn})	type (φ_{bn})	share (s_{bn})	appeal (A_{bn})	type (φ_{bn})
Beer						
brand market share	1.566					
brand appeal	0.694	1.397				
brand type	0.989	0.682	1.651			
firm market share	-0.571	-0.254	-0.537	0.545		
firm appeal	-0.368	-0.397	-0.312	0.683	0.298	
firm type	-0.557	-0.234	-0.526	0.993	0.656	0.515
Spirits						
brand market share	0.929					
brand appeal	0.718	0.801				
brand type	0.999	0.716	0.972			
firm market share	-0.463	-0.231	-0.465	0.369		
firm appeal	-0.367	-0.324	-0.376	0.731	0.164	
firm type	-0.451	-0.224	-0.454	0.996	0.737	0.401

Diagonal: ratio of FE variances to variance of the dependent variable. **Off-diagonal:** correlation. Underlying regressions keep the largest connected set.

C The cost of quality

We can back out measures of brand-level quality—defined as the destination- and time-invariant component of appeal—and marginal production cost. The method is to regress the inferred $\ln A_{bn}$ and $\ln c_{bn}$ on "friction" determinants and recover the brand fixed effects. An interesting question is whether "quality pays." There is already an empirical literature supporting the intuition that making higher quality brands requires higher production costs.[34] If the elasticity of production costs with respect to quality exceeds one, then higher quality brands will have lower market shares and, presumably, lower profits.[35] Figure 9 shows the relationship between brand fixed effects for appeal and cost for three categories under the Bertrand conduct assumption. Quality is expensive in the wine industry with an elasticity of 0.75, but relatively cheap for Bottled Water, with an estimated elasticity of 0.25. Beer, as well as Spirits and Carbonates (the last two shown in Figure 10 in the appendix) exhibit intermediate cost-quality elasticities.[36] Since all the estimated cost of quality elasticities are well below one, we infer that quality is profitable in these industries.

Figure 9: How higher quality affects costs

(a) Wine (b) Beer (c) Bottled Water

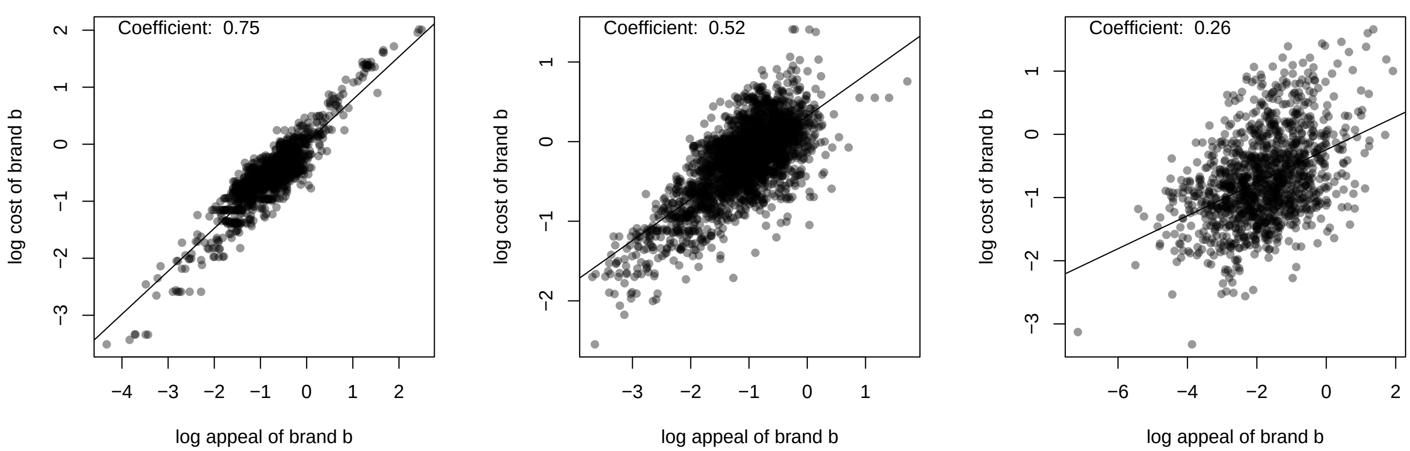

Each point is a brand fixed effect for appeal (Table 4, column 2) and its associated fixed effect for cost (Table 4, column 3).

[34]See, for example, Kugler and Verhoogen (2011) and Crozet et al. (2012).

[35]The presumably can be removed under monopolistic competition. With oligopoly, profits and sales are not strictly proportional but we conjecture they are still monotonically related.

[36]Sorted by the cost of quality elasticity, the estimates for Bertrand and Cournot (respectively) are generally very close to each other: Wine (0.70, 0.71), Spirits (0.52, 0.51), Beer (0.45, 0.37), Carbonates (0.26, 0.21), and Bottled Water (0.24, 0.21).

Figure 10: How higher quality affects costs (other products)

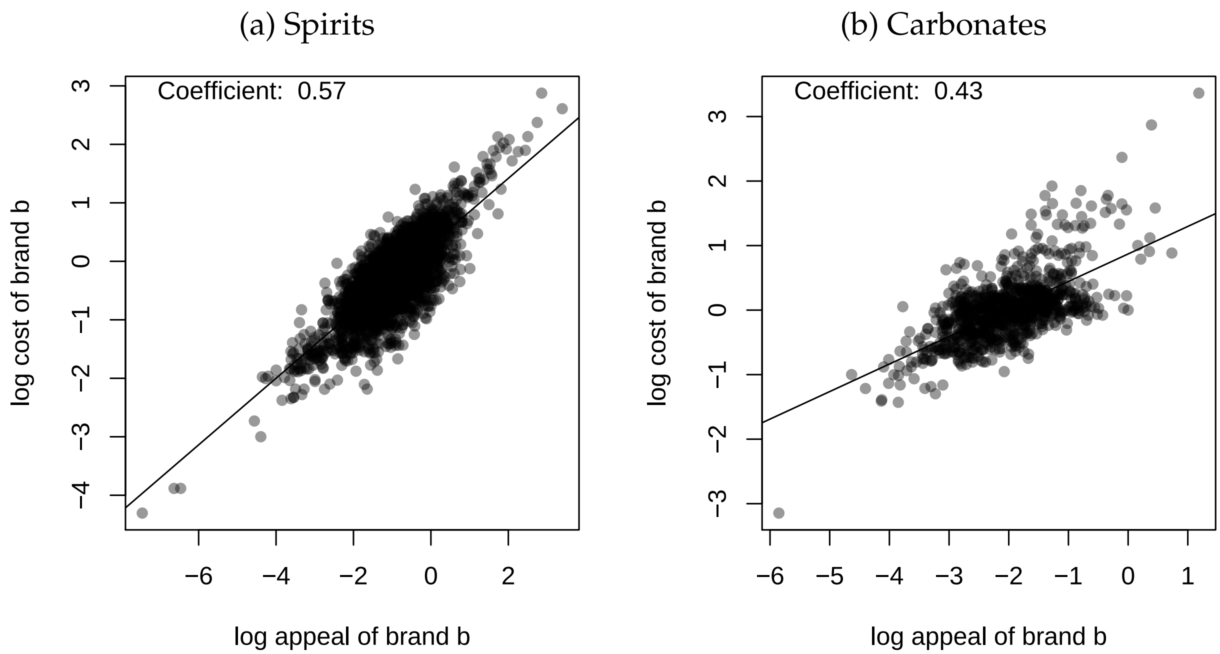

Each point is a brand fixed effect for appeal (Table 4, column 2) and its associated fixed effect for cost (Table 4, column 3).

D Ownership counterfactuals 2007–2018 (supplemental)

Figure 11: Effects of beer brand ownership changes in 2007–16

(a) Bertrand: pure market power effects ($\hat{\varphi} = 1$)

(b) Bertrand: including HQ changes ($\hat{\varphi} \neq 1$)

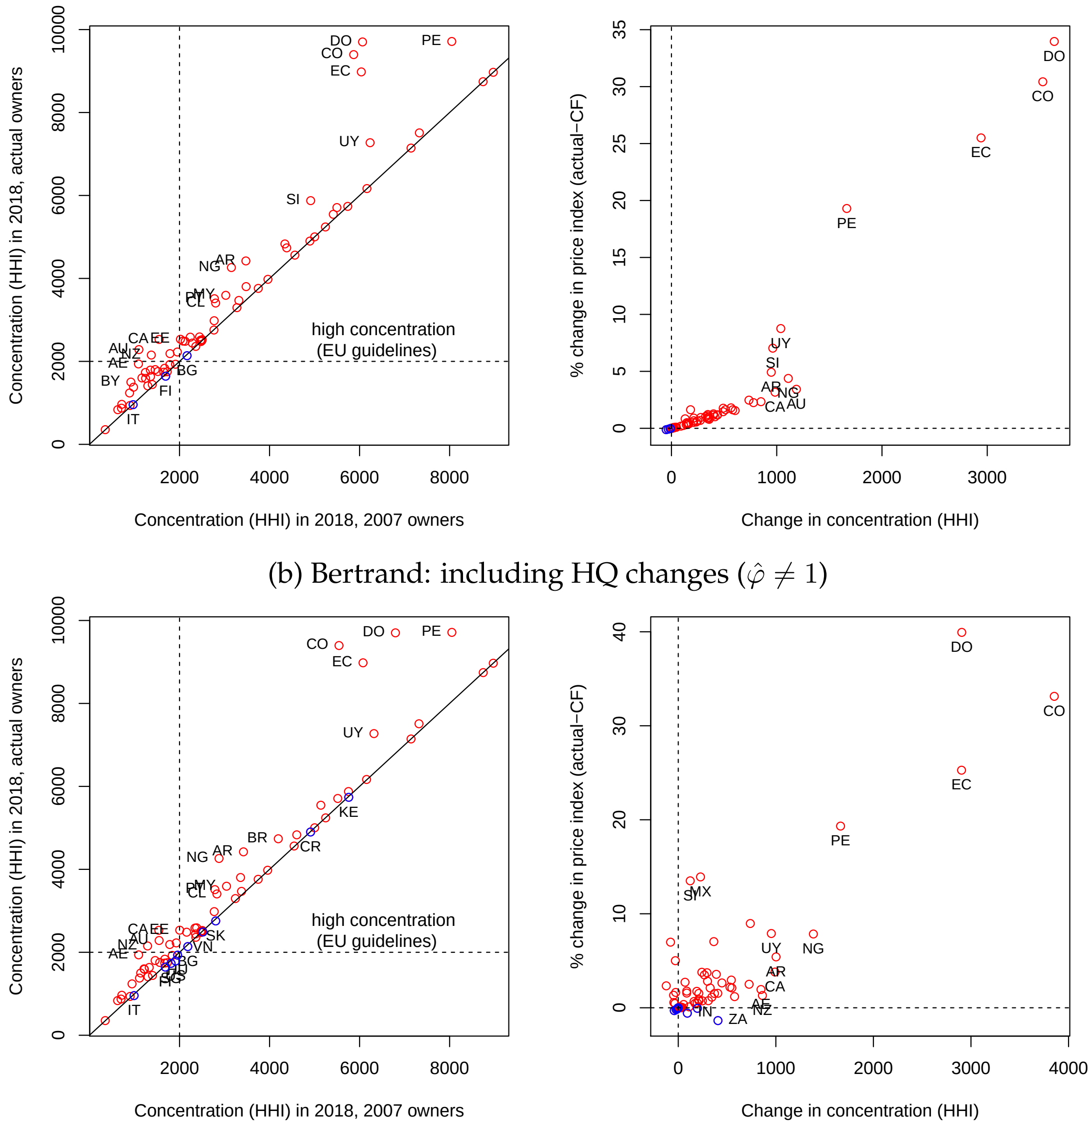

Figure 12: Effects of spirits brand ownership changes 2007–16

(a) Bertrand: pure market power effects ($\hat{\varphi} = 1$)

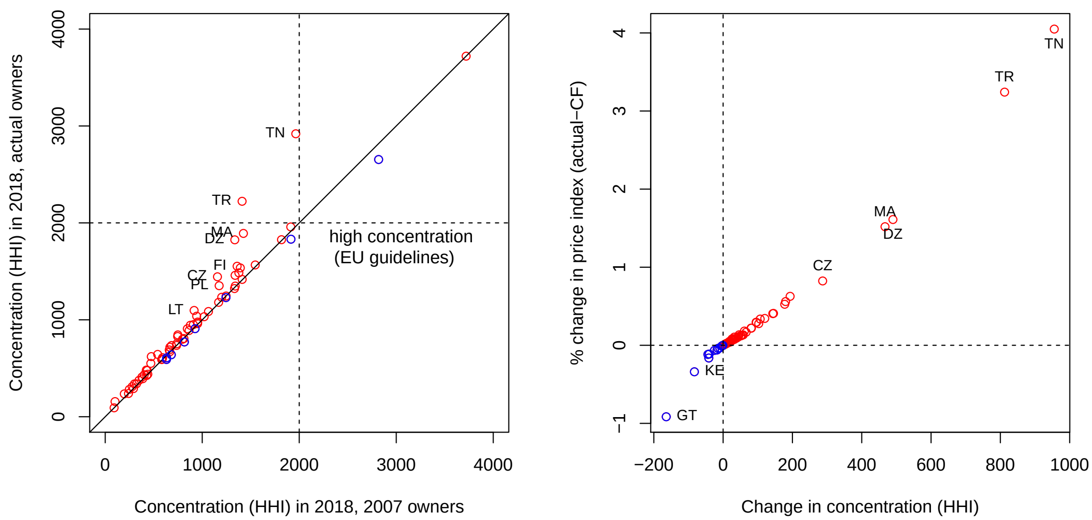

(b) Bertrand: including HQ changes ($\hat{\varphi} \neq 1$)

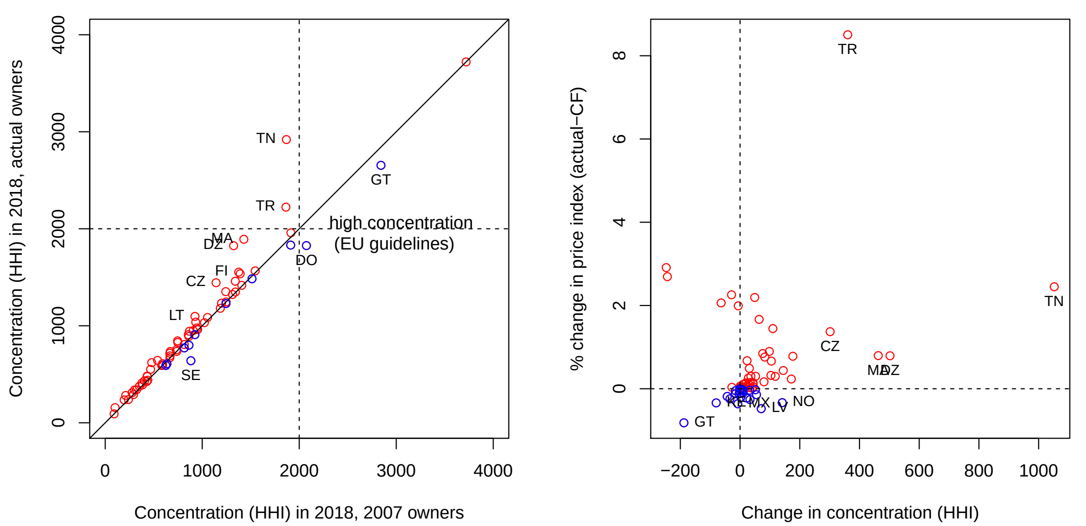

E Concentration and markups

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. We assume that there is a very large number of fringe firms, such that we can treat them as massless, and therefore assign them the monopolistically competitive Lerner index of $L_o = 1/\sigma$. The zero mass assumption implies that the Herfindhal index in market n is $H_n = \sum_{f \neq o} S_{fn}^2$.

There appears to be no consensus on the preferred way to specify and aggregate the markup. De Loecker and Eeckhout (2017) 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) 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

$$L_n^h \equiv \left(\sigma S_{on} + \sum_{f \neq o} \frac{S_{fn}}{L_{fn}} \right)^{-1} = \frac{1}{\sigma - (\sigma - \eta) H_n}. \tag{33}$$

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,

$$L_n^a \equiv \frac{1}{\sigma} S_{on} + \sum_{f \neq o} S_{fn} L_{fn} = \frac{1}{\sigma} + \left(\frac{1}{\eta} - \frac{1}{\sigma} \right) H_n \tag{34}$$

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

$$\mu_n^h \equiv \left(\frac{\sigma - 1}{\sigma} S_{on} + \sum_{f \neq o} \frac{S_{fn}}{\mu_{fn}} \right)^{-1} = \left[\frac{\sigma - 1}{\sigma} - \left(\frac{1}{\eta} - \frac{1}{\sigma} \right) H_n \right]^{-1} \tag{35}$$

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 $H_n \to 1$. The general point is that under both types of conduct, aggregate markups are increasing with the Herfindahl, moving from monopolistic competition to monopoly levels. Unfortunately, we have not found any single measure of the aggregate markup that has a closed form relationship with the Herfindahl under both Cournot and Bertrand conduct.

De Loecker and Eeckhout (2017) use a market share weighted price to cost ratio, that is

$$\mu_n^a \equiv \frac{\sigma}{\sigma - 1} S_{on} + \sum_{f \neq o} S_{fn} \mu_{fn} \approx 1 + L_n^a, \tag{36}$$

where the last approximation comes from the Hall (2018) approximation that $\mu_{fn} \approx 1 + L_{fn}$.

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