Concentration without Greater Market Power? Evidence from U.S. Sugar's Acquisition of Imperial Sugar

Chung-Yen Lo

Abstract

This paper examines whether a merger between sugar refiners affects competition in the retail markets where their brands compete. I study U.S. Sugar's acquisition of Imperial Sugar, announced in March 2021, using NielsenIQ Retail Scanner data from 2016 to 2022. A difference-in-differences design comparing the Southeast, where the merging firms' refining capacity is concentrated, with other regions finds no relative increase in retail prices following the announcement. I then estimate a mixed logit demand model and recover markups and marginal costs under multiproduct Bertrand competition. Substitution is spread across many competing brands rather than concentrated among the merging parties' products. Taking the ownership changes as known, a counterfactual decomposition separates the ownership and cost channels within a common demand environment. The change in ownership raises equilibrium prices by 0.03 percent and leaves markups, profits, consumer surplus, and welfare essentially unchanged, while differences between the pre- and post-announcement cost environments lower prices by 0.4 percent and markups by 2.1 percent. The merging firms together account for under 4 percent of retail volume, and private label, priced by retailers, accounts for most of the category, so the transaction alters the ownership of only a small share of the products consumers choose among. The results suggest that the competitive consequences of a merger among producers depend less on the change in producer concentration than on how much of the consumer-facing market the merging firms control and on how readily buyers substitute between the products they own.

1 Introduction

Sugar is a staple of household food expenditure in the United States, and the retail market for it is large. Annual per capita consumption of refined sugar exceeds 1,600 ounces, making the retail sugar market economically significant. The refining stage that supplies this market, however, is concentrated and geographically clustered: a small number of firms operate the country's major refineries, with substantial capacity in the Southeast and Gulf regions. When an industry combines concentrated production with a differentiated retail market, consolidation among producers raises the question of whether greater concentration at one stage translates into market power at the next.

On March 21, 2021, U.S. Sugar announced its acquisition of Imperial Sugar, including Imperial's Port Wentworth refinery in Georgia. The transaction combined two important participants in the U.S. sugar supply chain. The Department of Justice challenged the acquisition, arguing that concentration in the relevant regional market was already high and that the transaction would substantially lessen competition.[1] In July 2023, the U.S. Court of Appeals for the Third Circuit rejected the government's effort to block the transaction. The court emphasized the role of wholesalers and national retail distribution networks in allowing refined sugar to move across geographic markets, raising the possibility that geographic concentration in refining may not translate directly into local downstream market power.

This paper asks whether that possibility is borne out in retail outcomes. The question matters beyond this case. When a merger raises concentration among producers, its downstream consequences depend on two distinct channels. Common ownership changes the set of products over which a firm internalizes substitution, which raises pricing incentives mechanically. Changes in sourcing, production, or distribution can move marginal costs in either direction. The two channels can offset one another, so the empirical question is not whether concentration rose but whether the resulting configuration of ownership and costs

[1] United States v. United States Sugar Corp. For the complaint, see `https://www.justice.gov/atr/case-document/file/1451211/dl?inline`.

raised downstream price–cost margins.

Answering it is difficult because refiner-level sales, wholesale prices, supply contracts, and marginal costs are not publicly observed. Retail scanner data offer an indirect route. Major refiners either own consumer-facing brands or supply sugar sold under other brands, including private label, so an ownership change upstream can alter both the cost environment facing retail products and the pricing incentives of firms controlling multiple brands. Observing retail prices and quantities allows those consequences to be measured even when the underlying supply relationships are not.

I use NielsenIQ Retail Scanner data covering grocery, drug, and mass-merchandise outlets from 2016 to 2022. I define a market as a state–quarter and a product as a brand–sugar type combination, so that granulated, brown, and powdered sugar compete within a single retail sugar category. The Southeast receives particular attention because it contains the refining capacity most directly affected by the transaction.

The analysis proceeds in two stages. First, difference-in-differences and dynamic event-study specifications compare retail outcomes in the Southeast with those elsewhere around the announcement of the acquisition. Because the transaction did not close until November 2022, at the very end of the sample, these specifications measure the post-announcement period rather than a completed change in ownership; I interpret them accordingly. I also examine whether retail prices vary with geographic proximity to refining capacity.

Second, I estimate a random-coefficients logit demand model in the spirit of Berry et al. (1995) and Nevo (2001). Consumers differ in price sensitivity through a normally distributed random coefficient on price, which allows the products a consumer substitutes toward to depend on how price sensitive that consumer is. The model includes brand, sugar-type, state, and quarter fixed effects, together with sugar-type by quarter-of-year effects that absorb the pronounced seasonality of brown and powdered sugar, and it instruments for price using the leave-one-out price of the same product in other states, world raw sugar futures prices interacted with sugar type, and a differentiation instrument built from predicted prices. I

then combine the estimated demand derivatives with a multiproduct Bertrand pricing model to recover product-level markups and marginal costs. Throughout, brand owners are modeled as setting retail prices; I do not model refiner–retailer contracting separately, so recovered marginal costs should be read as the full cost behind the shelf price rather than as a refiner's production cost.

The reduced-form results provide little evidence of a price increase in the Southeast around the announcement. Prices show no clear relative increase, and the dynamic estimates reveal no discrete break at the announcement date. Retail prices do vary with proximity to refining capacity, but the estimated gradients are modest.

The structural results point the same way. Recovered marginal costs vary widely across products and markets, with store brands cheaper to supply than national brands. The distributions before and after the announcement overlap substantially, and neither markups nor Lerner indices display a common, persistent increase. Imperial and Dixie Crystals, the two acquired brands, follow distinct cost and pricing paths rather than a shared post-announcement expansion in margins.

Finally, I use the estimated model to separate the ownership channel from the cost channel. Evaluating four combinations of pre- and post-period ownership and marginal costs within a common demand environment, I find that the ownership change alone has a negligible effect on equilibrium prices, markups, profits, consumer surplus, and welfare. Replacing pre-period with post-period marginal costs while holding ownership fixed lowers equilibrium prices by about 5.6 percent and absolute markups by about 3.1 percent; consumer surplus rises by about 0.9 percent while total welfare moves by about 0.04 percent. The economically meaningful price response therefore comes from the cost environment, not from the change in ownership.

This decomposition also disciplines the interpretation of raw pre- and post-period comparisons. Market-level variable profits are higher and consumer surplus lower in the later period, but those changes combine shifts in market size, product composition, demand, costs,

and ownership. Holding demand fixed and varying only ownership isolates the channel that merger policy is concerned with, and that channel is close to zero here.

Two features of the setting explain why. The merging parties' combined retail presence is small: U.S. Sugar accounts for a negligible share of retail volume in these data, and Imperial for a modest one, so the additional substitution internalized by common ownership is limited. And the retail category is dominated by private label, whose pricing is not directly altered by the transaction. Production concentration, in this case, does not propagate mechanically into retail margins.

More broadly, the results suggest that evaluating consolidation in supply chains requires more than measuring the change in upstream concentration. Whether that change reaches consumers depends on how the merging firms' products substitute for one another downstream, how much of the downstream market they actually touch, and how costs move alongside ownership.

Roadmap. Section 2 reviews the related literature. Section 3 describes the industry and the transaction. Section 4 introduces the data and documents the main features of the retail sugar market. Section 5 presents the reduced-form evidence. Section 6 sets out the demand model and estimation. Section 7 reports demand estimates, markups, and marginal costs. Section 8 presents the counterfactual decomposition. Section 9 concludes.

2 Literature Review

This paper draws on three literatures: the structural evaluation of mergers in differentiated-product markets, the effects of mergers whose competitive overlap is measured in a different market from the one where the transaction occurs, and the industrial organization and history of the U.S. sugar industry. Its contribution is to the first and third; it relates to the second by asking whether a transaction between two refiners is reflected in the retail markets where their brands compete.

Berry (1994) and Berry et al. (1995) develop the discrete-choice and random-coefficients frameworks underlying demand estimation with endogenous prices, and Nevo (2000, 2001) shows how to apply them to consumer-product markets to recover elasticities, markups, and substitution patterns and to simulate alternative ownership structures. This approach evaluates mergers through product-level substitution and equilibrium pricing incentives rather than through concentration alone.

A related literature assesses how well such simulations predict realized outcomes. Peters (2006) compares predicted and realized price effects of airline mergers and shows that performance depends on the demand specification and on post-merger adjustment. Weinberg and Hosken (2013) find substantial variation across consumer-product mergers between simulated and observed effects. Miller and Weinberg (2017) study the MillerCoors joint venture and find price increases larger than a move between Nash–Bertrand equilibria would imply, using structural counterfactuals and institutional evidence to distinguish the mechanisms. Their analysis shows that ownership changes can alter incentives in ways that depend on the competitive environment.

The present setting differs from most applications in one respect. Structural merger studies typically examine firms that compete directly in the market where outcomes are measured. Here the transaction occurs among refiners while the outcomes are observed in retail markets containing many brands and extensive private label. The differentiated-product framework is used to ask whether the ownership change reaches downstream prices, markups, and welfare at all.

A long literature studies mergers and ownership changes whose competitive consequences appear in a market other than the one in which the transaction takes place. Theoretical work in this tradition examines successive markups (Spengler, 1950), foreclosure (Ordover et al., 1990), and contracting and commitment (Rey and Tirole, 2007). Empirically, Hastings (2004) and Hastings and Gilbert (2005) study how ownership structure relates to retail gasoline prices; Villas-Boas (2007) and Bonnet and Dubois (2010) estimate models of

manufacturer–retailer relationships in yogurt and bottled water; and Crawford et al. (2018) examine integration in multichannel television. What these studies share with the present setting is that a change in ownership at one stage of production coincides with a change in which firms control competing consumer products, and that the competitive question is settled in the market where those products meet.

The question I ask is narrower than the one these papers address, and the model is correspondingly simpler. I do not observe refiner-level production, wholesale prices, contract terms, or bargaining outcomes, and I do not model contracting between refiners and retailers. What the 2021 acquisition changes in my data is which retail brands are priced by the same firm. I therefore treat the transaction as a reassignment of brands across pricing entities and ask whether retail prices, substitution patterns, markups, and welfare respond. The supply model is a single-stage multiproduct Bertrand game among the firms that own retail brands, so recovered marginal costs summarize the cost of delivering a product to the shelf, inclusive of any margins earned earlier in the supply chain, and are not refiner production costs. Nothing in the analysis requires taking a position on how sugar moves from refineries to retailers, or on the terms under which it does.

Atalay et al. (2014) support treating the question this way. Using establishment-level data, they show that many upstream establishments within a firm ship little to that firm's own downstream establishments, so common ownership need not correspond to physical input flows. In this setting the relevant question is therefore which retail brands came under common control, not which refinery supplies which brand.

Genesove and Mullin (1998) are the closest precedent for the supply-side approach used here: they infer marginal costs and conduct in sugar refining from observed prices and demand information, showing what can be recovered when costs and conduct are unobserved. Genesove and Mullin (2001) analyze the Sugar Institute and document how information exchange and standardized business practices supported coordination without explicit price fixing, and Genesove and Mullin (2006) study price wars following entry. Eichner (1969)

provides the broader historical account of consolidation in refining, and Moss and Schmitz (2001) examine how trade policy and organization shape the distribution of rents along the sugar marketing chain. Beghin and Elobeid (2015) quantify the welfare effects of the U.S. sugar program.

This literature concerns the refining stage, the policy environment, and the historical organization of the industry. Comparatively little is known about the modern retail market, where products are differentiated by brand, sugar type, package size, and attributes such as organic certification, and where private label accounts for the majority of volume. This paper provides that downstream perspective: it estimates a random-coefficients demand system for retail sugar, characterizes substitution across brands and sugar types, and recovers markups and marginal costs at the product level.

The paper also speaks to the literature on concentration-based merger screening. Nocke and Whinston (2022) show that the competitive implications of a merger depend on the change in concentration and the underlying mechanism rather than on the post-merger level alone. The present case illustrates a complementary point: a transaction that raises concentration substantially at one stage may change the downstream ownership structure very little, because the merging firms' combined presence in the downstream market is small.

3 Industry Background

The U.S. sugar industry is characterized by substantial concentration at the refining stage and has long been shaped by federal policies governing domestic production and imports. The U.S. sugar program, administered by the U.S. Department of Agriculture (USDA), combines tariff-rate quotas with domestic price-support mechanisms, including marketing allotments and nonrecourse loans. These policies restrict the quantity of sugar that can enter the U.S. market at low tariff rates and support domestic sugar prices. As a result, U.S. sugar prices have historically remained above world prices, while domestic production

supplies the majority of U.S. sugar consumption. The combination of policy protection and a concentrated refining sector provides an important institutional setting for studying how changes in upstream ownership may affect downstream competition.

Government intervention in U.S. sugar markets dates to the early years of the republic and expanded substantially during the nineteenth and twentieth centuries. The industry subsequently experienced a long process of concentration in refining. By the early twentieth century, the American Sugar Refining Company controlled a large share of U.S. refining capacity. Following the company's dissolution under the Sherman Act, the refining industry evolved toward an oligopolistic structure characterized by a small number of major firms. Historical studies of the industry document how concentration, strategic interaction, and institutional arrangements shaped competition in sugar refining (Eichner, 1969; Genesove and Mullin, 1998, 2001, 2006). This historical background is relevant for the present study because the refining sector remains concentrated even though the retail market contains a much larger number of differentiated products.

Refining capacity is also geographically concentrated. Sugarcane production is concentrated primarily in Florida and Louisiana, while sugar beet production is concentrated in the Midwest and Western regions of the United States (Moss and Schmitz, 2001). Refining facilities are consequently located unevenly across the country, reflecting both the geographic distribution of raw sugar production and the economics of transportation. Although refined sugar can be transported over substantial distances, transportation and distribution costs create geographic differences in the cost of supplying downstream markets. Geographic proximity to refining facilities may therefore influence the competitive environment faced by retail brands and provide an important source of regional variation.

Figure 1 illustrates the geographic distribution of major refineries. Refining facilities associated with United Sugars Cooperative are concentrated across the Midwest and Gulf regions, while U.S. Sugar operates refining capacity in Florida. Imperial Sugar operated the Port Wentworth refinery in Georgia and marketed refined sugar under the Imperial

and Dixie Crystals brands. The geographic separation of these facilities is important for understanding the regional exposure to the 2021 acquisition and motivates the particular focus on the Southeast in the reduced-form analysis.

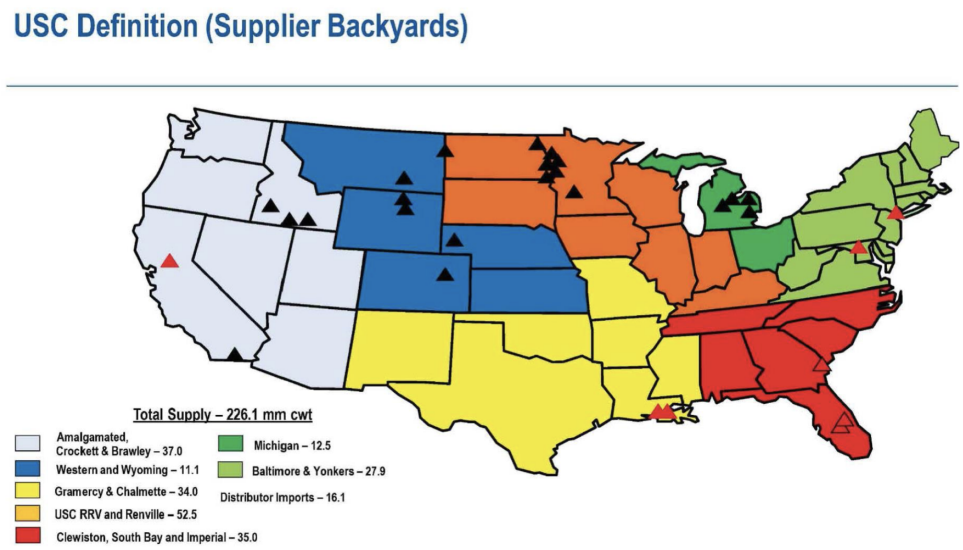


Figure 1: Geographic Distribution of Sugar Refineries

The 2021 transaction is of particular interest because the two firms occupied different positions in the retail sugar market. Imperial Sugar operated the Port Wentworth refinery in Georgia and competed directly for consumers through its own brands, principally Imperial and Dixie Crystals, which are sold in grocery and mass-merchandise outlets across much of the eastern United States. U.S. Sugar, by contrast, has no comparable consumer-facing brand of its own. It reaches retail shelves through the United Sugars cooperative, which markets and sets prices for the sugar produced by its member refiners. The acquisition therefore did not combine two firms selling closely competing branded products. It brought an established branded portfolio under the same pricing umbrella as a cooperative whose retail presence takes a different form, and it did so while leaving the underlying refining assets in place.

The transaction can affect retail outcomes through two channels, and the two need not work in the same direction. The first operates through pricing incentives. When products

that were previously priced by separate firms come under common control, the firm setting prices internalizes substitution between them, which raises the incentive to price higher on both. The strength of this channel depends on how readily consumers substitute between the affected products: if a price increase by one diverts demand mainly to products the firm does not own, little is internalized and the incentive to raise prices is correspondingly weak. The second channel operates through costs. Changes in sourcing, refining, packaging, or distribution can move the cost of supplying retail products in either direction, and any such change is passed through to prices to a degree determined by the same demand system. Because the first channel tends to raise prices and the second may lower them, the net effect on retail prices is an empirical question rather than one that can be settled by observing the change in concentration.

Two institutional details determine how the transaction maps into the ownership structure used in the empirical model. The first concerns U.S. Sugar. According to the Department of Justice complaint, U.S. Sugar sells its refined sugar through United Sugars rather than marketing it independently, and the cooperative sets prices for the sugar produced by all of its members. The pricing incentives attached to U.S. Sugar's retail volume are therefore those of the cooperative rather than of U.S. Sugar acting alone, and I treat the members of United Sugars as a single pricing entity throughout the supply model. The second concerns Imperial. Its brands competed at retail as an independent portfolio before the transaction, and its sugar was to be marketed through United Sugars once the acquisition closed. The relevant change in ownership for retail competition is therefore the transfer of the Imperial and Dixie Crystals brands into the cooperative's pricing decisions, and it is this reassignment, rather than the transfer of refining capacity, that the counterfactual analysis evaluates.

Although refined sugar is nearly homogeneous as a physical commodity, retail sugar products are differentiated in ways that matter for consumer choice. Products differ by brand, sugar type, package size, and attributes such as organic certification, and the market

contains both national brands and private label. Prices and shares vary substantially across products and markets as a result, which motivates a differentiated-product demand model rather than a homogeneous-commodity framework. Because a brand may offer granulated, brown, and powdered products, substitution can occur both across brands within a type and across types; the demand model accommodates both and allows responses to be aggregated to the brand or category level.

These features shape the empirical design. The ownership change occurred upstream, where capacity is geographically concentrated, while the outcomes studied here are downstream. Whether the first reaches the second depends on substitution among retail products, on pass-through of any cost changes, and on the availability of alternative supply. Wholesalers, national chains, and dispersed refining capacity all allow retailers to source from outside a region, which is precisely the argument the courts found persuasive. The Southeast offers a useful testing ground: it contains the affected refining capacity while its retail markets contain many competing national and private-label products.

4 Data and Suggestive Evidence

4.1 Data

I use data from the NielsenIQ Retail Scanner dataset, made available through the Kilts Center for Marketing at the University of Chicago Booth School of Business. The dataset records weekly transactions from a broad panel of grocery, drugstore, and mass-merchandise chains across the United States and has been widely used in empirical industrial organization research (e.g., Hendel and Nevo (2006); Dubé et al. (2012); Miller and Weinberg (2017); Conlon and Rao (2023)).

The scanner data provide detailed product-level information, including Universal Product Codes (UPCs), brand names, manufacturer identifiers, product descriptions, and package sizes. For each UPC–store–week observation, the dataset reports total units sold, dollar sales,

and promotional indicators. I use these variables to construct prices per ounce and total quantities sold in ounces, allowing prices and quantities to be measured consistently across products with different package sizes. The data therefore provide detailed information on both product characteristics and consumer purchasing behavior at a highly disaggregated level. I focus on refined sugar products sold between 2016 and 2022, which provides five years of observations before the acquisition was announced in March 2021 and seven quarters afterwards. The analysis focuses on three major product types: granulated, brown, and powdered sugar. Within these categories, products differ along additional dimensions, including brand, sugar source, and organic certification. I exclude specialty sweeteners and sugar substitutes in order to maintain a consistent product-market definition. The final descriptive sample contains 97,828 brand–state–quarter observations.

Because the merger occurred at the refining stage and refining capacity is geographically concentrated, geographic market definition is central to the analysis. I define markets at the state–quarter level. State-level aggregation preserves meaningful geographic variation in refinery locations, transportation costs, and exposure to the merging firms, while quarterly aggregation reduces short-run variation arising from temporary promotions and other high-frequency retail shocks. Weekly UPC-level observations are therefore aggregated to the brand–state–quarter level using total ounces sold as weights when constructing prices. A product is defined as a brand paired with a sugar type, so a brand offering both granulated and brown sugar contributes two products to a market. The typical state–quarter contains about 71 products supplied by 39 brands.

For each brand–state–quarter observation, I compute the average price per ounce by dividing total dollar sales by total ounces sold and adjusting for inflation. Market shares are constructed using volume sold relative to the corresponding measure of market size in each state–quarter. In the structural demand analysis, I also construct an outside option representing purchases not accounted for by the observed inside goods, consistent with the discrete-choice framework developed below.

The descriptive and reduced-form analyses use the full sample of 97,828 observations. The structural demand estimation imposes additional requirements that market shares, outside shares, prices, product characteristics, and instruments be positive, finite, and observed, and that each product be sold in at least two states within a quarter so that the leave-one-out price instrument can be constructed. I also restrict attention to prices between 2 and 50 cents per ounce, which removes specialty items such as maple, vanilla, and cocktail sugars together with a small number of implausible unit values; these observations account for roughly 4.5 percent of the sample but 0.1 percent of volume. These restrictions reduce the demand-estimation sample to 88,270 observations. I retain the broader sample for the descriptive and reduced-form analyses because observations that remain informative about market outcomes need not be discarded solely because an instrument required for the structural estimation cannot be constructed.

The data reveal substantial heterogeneity across brands, sugar types, and geographic markets. Although a large number of brands appear in the national sample, sales are concentrated among a much smaller set of national brands and private-label products. Meaningful price dispersion also persists across brands and product types, and prices differ systematically across the three sugar types: granulated sugar, which accounts for about 79 percent of category volume, sells for 3.97 cents per ounce on a volume-weighted basis, while brown and powdered sugar sell for 6.62 and 6.70 cents. These patterns indicate that retail sugar exhibits economically relevant product differentiation despite the relative homogeneity of refined sugar as a physical product.

Two features of the measurement warrant emphasis for the structural analysis. First, because prices are quarterly unit values averaged across stores within a state, short-run promotional variation is averaged out. This attenuates measured price sensitivity relative to what store–week data would deliver, so the demand estimates below are best interpreted as a lower bound on price responsiveness and the implied markups as an upper bound. Second, the scanner data provide rich information on retail transactions but do not report whole-

sale prices, input costs, or contractual terms between refiners and retailers. Consequently, observed retail prices reflect equilibrium outcomes generated by interactions along the supply chain. In the structural analysis, I use estimated demand elasticities together with an assumed model of Bertrand pricing among brand owners to recover implied markups and marginal costs. This approach allows me to examine whether changes in retail prices reflect changes in marginal costs or changes in price–cost margins despite the absence of directly observed cost data, with the recovered costs measuring the full cost of delivering a product to the shelf rather than a refiner's production cost.

Overall, the NielsenIQ data provide broad coverage of retail sugar sales across the United States and allow me to combine descriptive evidence, reduced-form regional comparisons, and structural analysis to study how consolidation at the refining stage is reflected in retail outcomes.

4.2 Suggestive Evidence

Table 1 summarizes the retail sugar category. Panel A pools all state–quarter markets over the 2016–2022 sample period, so the brand counts report every brand observed anywhere in the United States at any point, rather than the number competing in a given market at a given time. On that basis product variety is extensive: the sample contains 327 granulated, 309 brown, and 192 powdered sugar brands, of which 177, 165, and 157 respectively are private label. Private label accounts for 60.7, 52.5, and 58.3 percent of volume in the three categories. Granulated sugar is the category staple, representing 79.2 percent of volume, with brown at 13.9 percent and powdered at 6.8 percent. Prices differ systematically across types. Granulated sugar is the cheapest at 3.97 cents per ounce on a volume-weighted basis, while brown and powdered sugar sell for 6.62 and 6.70 cents per ounce.

Local markets (Panel B) are considerably smaller than the national aggregates. A typical state–quarter contains about 31 granulated, 25 brown, and 15 powdered sugar brands, with private label accounting for 57 to 64 percent of volume. Sales within these markets are

concentrated: the brand-level Herfindahl–Hirschman index (HHI), computed over inside goods, averages 2,722 nationally and 2,972 in the Southeast, and the four largest brands account for 82 percent of inside-good volume in a typical market nationally and 84 percent in the Southeast. Within-type concentration is similar, averaging 2,750 for granulated, 2,690 for brown, and 3,171 for powdered sugar.

High concentration coexists with a long list of small brands, and the two facts are easier to reconcile once the arithmetic behind the index is made explicit. The HHI is a share-weighted average of brand shares, so an index of 2,722 says that the average ounce of sugar sold comes from a brand holding 27 percent of its local market. The median brand, by contrast, holds 0.09 percent. Both statements describe the same market: the first weights brands by the volume they sell, the second treats every brand on the shelf equally. In the median state–quarter, 33 brands compete, but only four hold at least five percent of volume, the largest alone holds 42 percent, and the 25 smallest hold under one percent each and together account for less than four percent of the category. Concentration therefore reflects the position of a handful of large brands, principally the leading private labels and national brands, while the long tail of specialty products contributes to variety without contributing to sales. Figure 2 shows the resulting distribution.

Table 1: Summary Statistics

Panel A: National, full sample

	Granulated	Brown	Powdered
# Brands	327	309	192
# Private-label brands	177	165	157
% Private-label volume	60.7	52.5	58.3
% Category volume	79.2	13.9	6.8
Price (cents/oz, volume weighted)	3.97	6.62	6.70
Price (cents/oz, median)	5.70	7.15	6.50

Panel B: Mean across state–quarter markets

	Granulated	Brown	Powdered
# Brands per market	31.3	24.7	15.3
% Private-label volume	64.0	56.9	58.8
HHI within sugar type	2750	2690	3171

Panel C: Market aggregates

	All states	Southeast
Products per market	71.3	62.4
Brands per market	38.9	34.5
Brand HHI (inside goods)	2722	2691
Top-4 brand share (%)	82.3	82.8

Notes: Panel A covers the full descriptive sample of 97,828 brand–state–quarter observations, 2016Q1–2022Q4. Panels B and C report means across state–quarter markets. Prices are deflated dollar sales divided by ounces sold. The Southeast comprises AL, FL, GA, NC, SC, and TN.

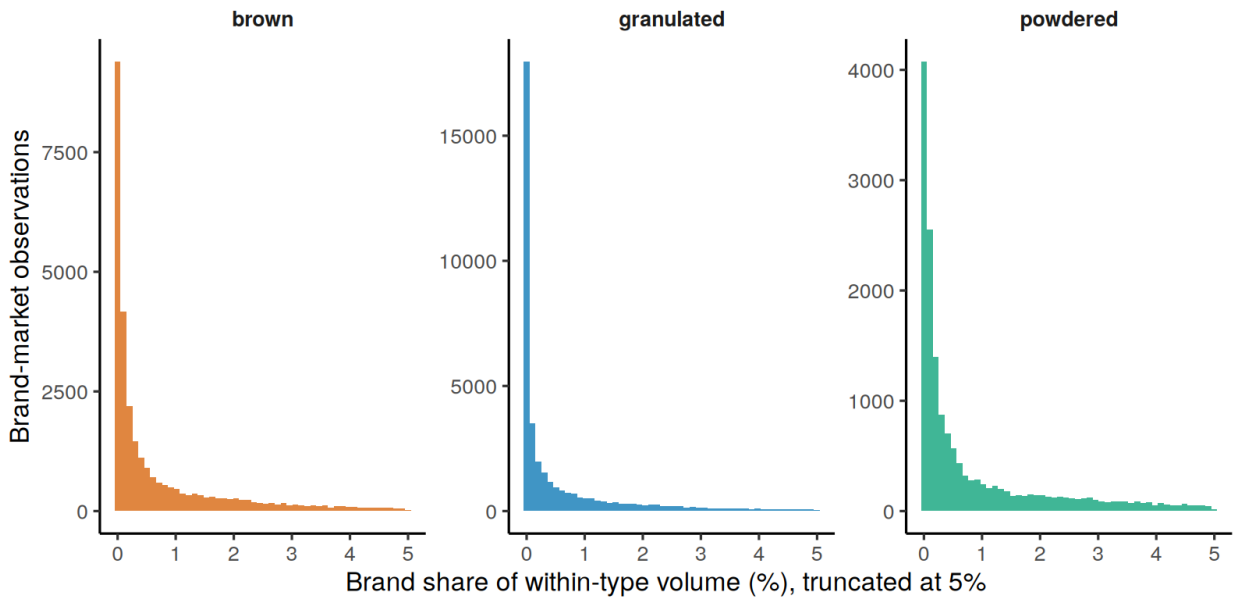


Figure 2: Distribution of brand market shares

Figure 3 shows how category composition evolves. Granulated sugar dominates throughout, accounting for roughly 80 percent of volume, though its share drifts down over the sample as brown sugar's rises. Composition also shifts from 2020 onward, coinciding with the increase in at-home food consumption during the COVID-19 pandemic, which underscores the need to control for aggregate time shocks when assessing changes around the announcement.

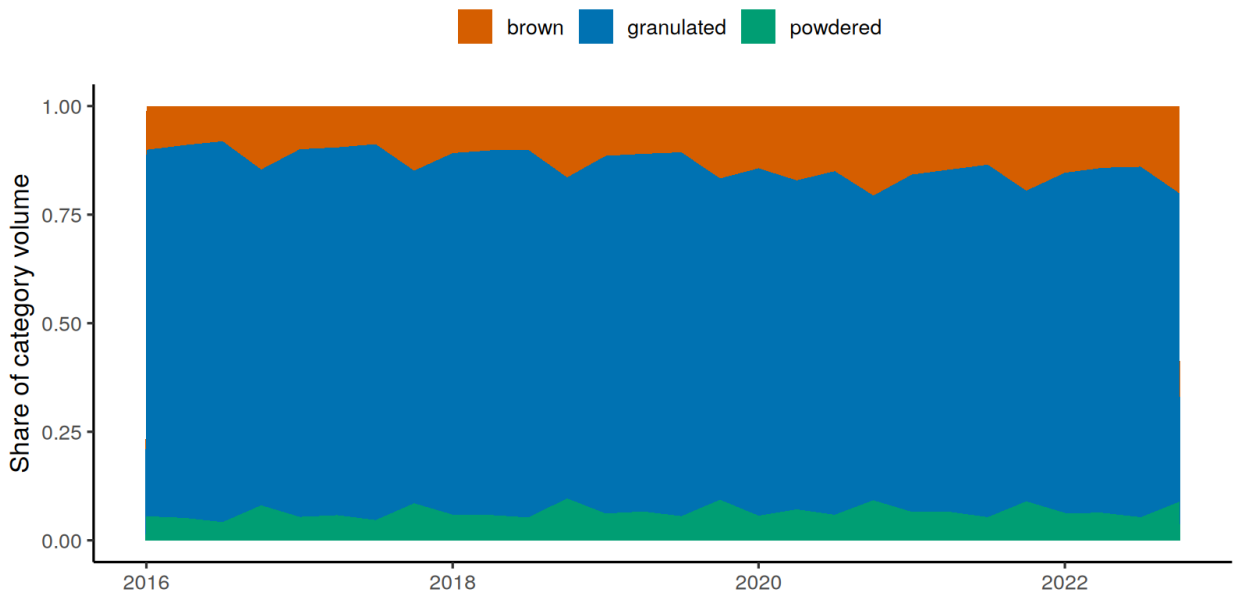


Figure 3: Quarterly share trends by sugar type

The figure also displays a pronounced sawtooth pattern, which Figure 4 quantifies. Relative to each year's average volume, brown and powdered sugar rise by 69 and 74 percent in the fourth quarter, against 23 percent for granulated. Holiday baking therefore shifts the composition of the category and not merely its level, so quarter fixed effects common to all products cannot absorb it. This motivates the sugar type by quarter-of-year fixed effects used in the demand specification; without them, the cost instruments interacted with sugar type fail their overidentification tests, since the seasonal component of type-specific demand is correlated with the seasonal component of the instruments.

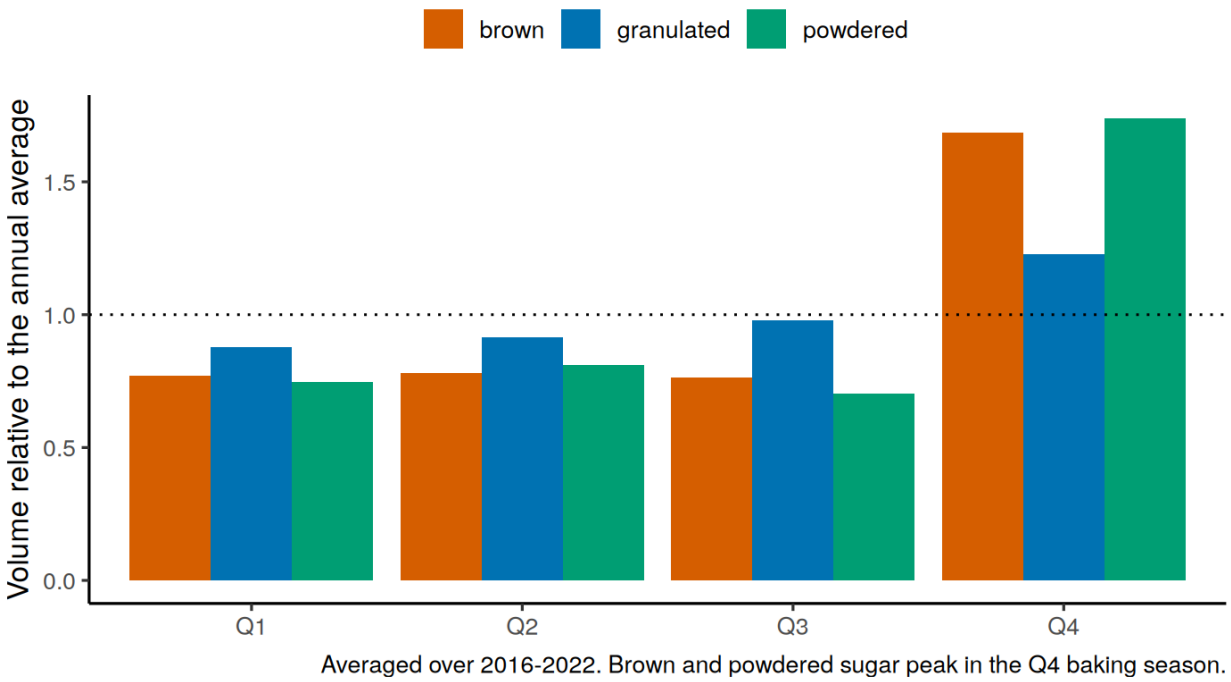


Figure 4: Seasonality of volume by sugar type, relative to the annual average

A final feature of the industry is relevant for interpreting the transaction. Several refiners reach retail markets both through their own brands and by supplying sugar sold under other brands, including private label. Changes in ownership or in the cost of refining may therefore affect a wider set of retail products than the merging firms' own brands. At the same time, the merging parties' direct retail presence is limited: Imperial and Dixie Crystals together account for 3.3 percent of national volume and 6.7 percent in the Southeast, while U.S. Sugar's retail brands account for less than 0.5 percent. Any ownership effect operating through common control of retail brands is therefore bounded by a small share of the category, a point the structural results confirm.

Taken together, the descriptive evidence establishes four facts that shape the analysis. Retail sugar offers extensive product variety but concentrated sales; private label accounts for the majority of volume; composition varies systematically across types, markets, and seasons; and the merging firms' combined retail footprint is small. These features motivate both the regional comparisons that follow and the structural model used to separate ownership from

cost effects.

4.3 Reduced-Form Evidence on Prices and Market Shares

As a first step, I begin with reduced-form evidence on how retail prices and quantities evolved around the announcement of the acquisition. These estimates ask a transparent question: did outcomes in the Southeast, the region most exposed to the transaction, move differentially relative to the rest of the country?

I estimate

$$Y_{bst} = \beta_1(Post_t \times Southeast_s) + \phi_b + \tau_{c(b)} + \eta_s + \lambda_t + \varepsilon_{bst},$$

where Y_{bst} is the log retail price or log volume of brand b in state s and quarter t, and ϕ_b, $\tau_{c(b)}$, η_s, and λ_t are brand, sugar-type, state, and quarter fixed effects. I define the post-merger period as beginning in 2021Q2 and define the Southeast as the set of core Southeastern states used throughout the baseline analysis.[2] State fixed effects absorb persistent differences across states, while quarter fixed effects absorb aggregate shocks common to all geographic markets. Brand and sugar type fixed effects further account for persistent differences across products. The coefficient β_1 measures the differential change in the Southeast after 2021Q2 relative to the contemporaneous change elsewhere.

To examine the dynamics of price adjustment and to assess the plausibility of the parallel-trends assumption, I first estimate an event-study specification that interacts the Southeast indicator with quarter dummies,

$$Y_{bst} = \sum_{\tau \neq 2020Q4} \beta_\tau \left(\mathbf{1}\{t = \tau\} \times Southeast_s\right) + \phi_b + \tau_{c(b)} + \eta_s + \lambda_t + \varepsilon_{bst},$$

where the omitted category is 2020Q4. Each β_τ measures the difference between the South-

[2]Core Southeastern states contain Alabama, Florida, Georgia, Mississippi, North Carolina, South Carolina and Tennessee.

east and other regions in quarter τ relative to that base quarter, after absorbing brand, sugar-type, state, and quarter fixed effects. Standard errors are clustered at the state level. Figure 5 plots the resulting coefficients and 95 percent confidence intervals for log prices.

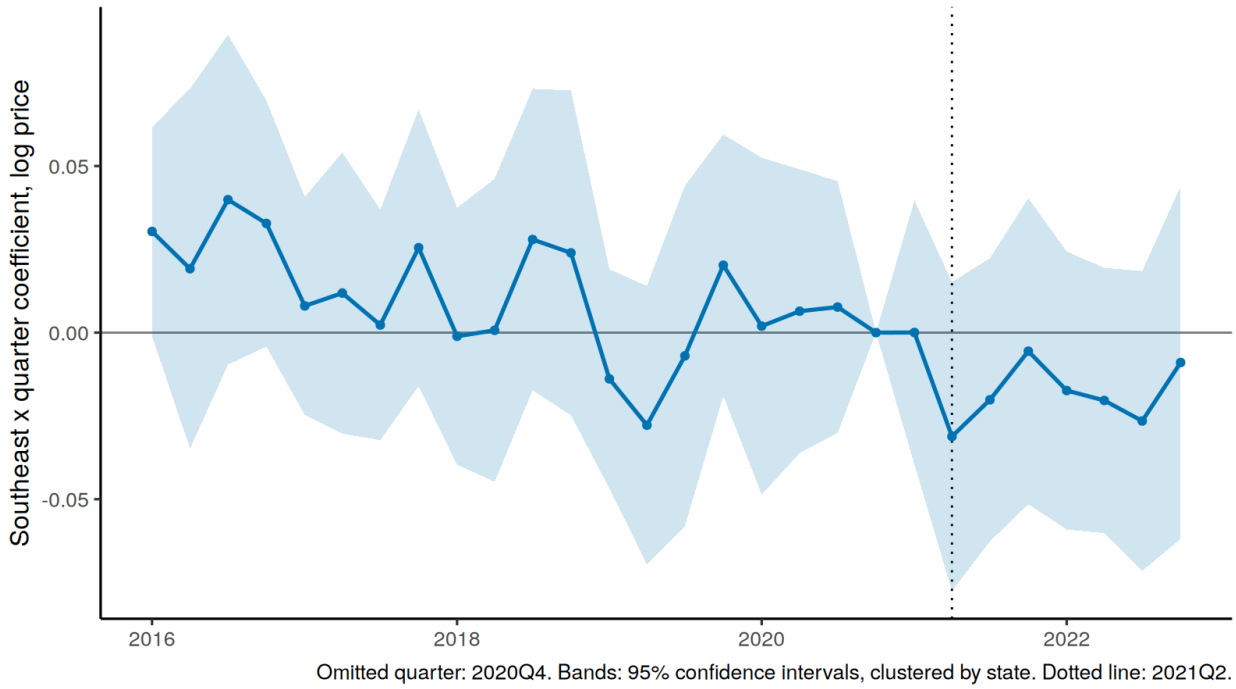


Figure 5: Event-study estimates for log price

The pre-announcement coefficients provide no evidence of differential trends. Across the twenty quarters preceding the base period they average 0.015 in absolute value, none is statistically distinguishable from zero at conventional levels, and the series fluctuates around zero without a discernible trend. The Southeast and the rest of the country were therefore on comparable price paths before the transaction, which is the identifying assumption the difference-in-differences design requires.

After 2021Q2 the coefficients remain small and statistically insignificant. They range from -0.03 in 2021Q2 to -0.005 in 2021Q4 and average -0.02 across the seven post-announcement quarters, with confidence intervals that contain zero in every quarter. There is no discrete break at the announcement and no subsequent drift: prices in the Southeast neither rose nor fell appreciably relative to prices elsewhere. The event-study evidence therefore provides

little support for a discrete increase in regional retail prices following the acquisition.

Table 2 reports the static estimates. The price coefficient is −0.028 and statistically significant at the one percent level, implying that Southeast prices rose about 2.8 percent less than prices elsewhere over the post-announcement period. The volume coefficient is 0.038 and is not statistically distinguishable from zero. Restricting the sample to granulated sugar, the category's dominant product and the one unaffected by the unit correction, gives a price coefficient of 0.008 (0.010), again indistinguishable from zero.

Table 2: Difference-in-Differences Estimates

	(1)	(2)	(3)	(4)
	log(Price)	log(Volume)	log(Price)	log(Volume)
	All types		Granulated only	
Post × Southeast	-0.028***	0.038	0.008	-0.108
	(0.01)	(0.059)	(0.01)	(0.08)
Brand FE	Yes	Yes	Yes	Yes
Sugar Type FE	Yes	Yes	No	No
State FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Observations	97,828	97,828	42,953	42,953

Notes: The post period begins in 2021Q2, the quarter after the acquisition was announced. The Southeast is AL, FL, GA, MS, NC, SC, and TN. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

These are relative, not absolute, changes: the negative price coefficient means that Southeast prices rose less than prices elsewhere, not that they fell. That distinction matters in a period when food prices were rising throughout the United States. Volume-weighted retail sugar prices in the sample rise from 4.48 to 4.81 cents per ounce between the two periods. Again, the reduced-form evidence provides little indication that the merger generated higher retail sugar prices in the region most directly exposed to the transaction. Instead,

the Southeast experienced a relative decline in prices.

The limitations of the reduced-form design motivate the structural analysis that follows. A null result on prices is not by itself a null result on market power. Retail prices reflect both margins and costs, and a price that does not rise is consistent either with unchanged pricing incentives or with an increase in margins offset by a decline in costs. Distinguishing the two requires separating price into a marginal cost and a markup, which the reduced-form design cannot do because costs are not observed. The structural model does exactly this: demand estimation recovers the substitution patterns that determine how much competitive pressure the merging firms' products exert on one another, inverting the pricing first-order conditions decomposes observed prices into markups and marginal costs, and simulating equilibrium prices under alternative ownership structures within a fixed demand environment isolates the effect of the ownership change from everything else that moved between 2016 and 2022.

4.3.1 Robustness to Alternative Treatment Definitions

Table 3 re-estimates the specification under narrower and broader definitions of the treated region. The price coefficient is -0.028 under the core definition, -0.017 under the narrow definition, and -0.019 under the broad definition; only the core estimate is significant at conventional levels. The volume coefficients are positive in all three cases and significant only under the broad definition. Still, there is no evidence of a relative price increase in the Southeast following the announcement.

Table 3: Robustness to Alternative Southeast Definitions

	log(Price)		
	Core	Narrow	Broad
Post × Southeast	-0.028***	-0.017*	-0.019
	(0.01)	(0.0100)	(0.012)
	log(Volume)		
	Core	Narrow	Broad
Post × Southeast	0.038	0.054	0.093**
	(0.059)	(0.086)	(0.048)
Brand, type, state, quarter FE	Yes	Yes	Yes
Observations	97,828	97,828	97,828

Notes: Standard errors clustered by state. The core region is AL, FL, GA, MS, NC, SC, TN; the narrow region is GA, SC, FL; the broad region adds KY, LA, VA, WV, and AR. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

5 Model

5.1 Setup

I observe the U.S. retail sugar market quarterly from 2016 to 2022. Markets are indexed by (s, t) for state s and quarter t. Products are indexed by j and defined by a brand $b(j)$ and a sugar type $c(j) \in \{\text{granulated}, \text{brown}, \text{powdered}\}$. Each brand is owned by a firm $f(b)$; private-label brands are assigned to the retailer that sells them, and the brands marketed through the United Sugars cooperative are assigned to a single pricing entity, since the cooperative sets prices for its member refiners' sugar.

Let $\mathcal{J}_{st}$ denote the products with positive sales in market (s, t). For each I observe the

retail price p_{jst}, quantity q_{jst} in ounces, and revenue. Market shares are $s_{jst} = q_{jst}/M_{st}$, where M_{st} is potential market size in ounces, and the outside share is $s_{0,st} = 1 - \sum_j s_{jst}$. The outside option represents sugar purchased outside the observed products and channels, and non-purchase.

Throughout, brand owners are modeled as setting retail prices. I do not model refiner–retailer contracting separately, so the marginal costs recovered are the full cost of delivering a product to the shelf, inclusive of any margins earned between the refinery and the register.

5.2 Demand

I estimate a random-coefficients logit model following Berry et al. (1995) and Nevo (2001). The indirect utility of consumer i from product j in market (s, t) is

$$U_{ijst} = \delta_{jst} + (\alpha_i - \bar{\alpha})\, p_{jst} + \varepsilon_{ijst}, \tag{1}$$

where δ_{jst} is the mean utility common to all consumers, $(\alpha_i - \bar{\alpha})p_{jst}$ captures the deviation of consumer i's price sensitivity from the population mean, and ε_{ijst} is an idiosyncratic taste shock distributed Type I Extreme Value. The utility of the outside option is normalized to $U_{i0st} = \varepsilon_{i0st}$, so that all utilities are measured relative to not purchasing sugar from the observed set of products.

Mean utility is specified as

$$\delta_{jst} = \bar{\alpha} p_{jst} + \beta x_{jst} + \gamma_{b(j)} + \tau_{c(j)} + \eta_s + \lambda_t + \kappa_{c(j),q(t)} + \xi_{jst}, \tag{2}$$

where p_{jst} is the retail price in cents per ounce, x_{jst} is the USDA organic seal rate, and ξ_{jst} is an unobserved product–market demand shock. The specification includes brand fixed effects $\gamma_{b(j)}$, sugar-type fixed effects $\tau_{c(j)}$, state fixed effects η_s, and quarter fixed effects λ_t. Brand fixed effects absorb persistent differences in quality, reputation, and packaging across the 297 brands in the estimation sample, including the private-label brands of individual retailers.

Sugar-type fixed effects absorb the substantial average differences across granulated, brown, and powdered sugar documented in Section 4.2, and state fixed effects absorb persistent geographic differences in the level of demand. Quarter fixed effects absorb national shocks common to all products, including input-cost movements, general food-price inflation, and the shift toward at-home consumption during 2020.

The term $\kappa_{c,q}$ denotes sugar-type by quarter-of-year fixed effects and merits separate comment. Section 4.2 documents that seasonality in this category is type-specific rather than common: relative to their annual averages, brown and powdered sugar volumes rise by 69 and 74 percent in the fourth quarter, against 23 percent for granulated. Holiday baking therefore shifts the composition of the category and not only its level, which quarter fixed effects common to all products cannot absorb.

Consumers differ in price sensitivity through a normally distributed random coefficient,

$$\alpha_i = \bar{\alpha} - \sigma\nu_i, \qquad \nu_i \sim \mathcal{N}(0,1), \tag{3}$$

where $\bar{\alpha}$ is the mean price coefficient and σ governs the dispersion of price sensitivity across consumers. This is the standard specification in the literature, and it is what distinguishes the model from the logit and nested logit alternatives reported alongside it in Section 7. Under logit, substitution between any two products is proportional to their market shares, so a price increase by one product diverts demand to its rivals in proportion to their size regardless of whether they are close substitutes. Allowing $\sigma > 0$ breaks that restriction. The consumers most likely to leave a product when its price rises are those with the strongest price sensitivity, and those consumers are disproportionately likely to switch to lower-priced alternatives. Substitution therefore concentrates among products that appeal to similar consumers rather than being spread in proportion to shares, which is what allows the model to generate realistic diversion between national brands and private label. The parameter σ is identified by variation across markets in the set of competing products and their prices:

when a product faces closer substitutes in price space, its demand is more elastic than the logit would imply, and the extent of that difference is informative about the dispersion of price sensitivity.

Given equation (3), market shares follow from integrating the logit choice probabilities over the distribution of ν:

$$s_{jst} = \int \frac{\exp\left[\delta_{jst} - \sigma\nu p_{jst}\right]}{1 + \sum_{k \in \mathcal{J}_{st}} \exp\left[\delta_{kst} - \sigma\nu p_{kst}\right]} \, d\Phi(\nu), \tag{4}$$

where Φ is the standard normal distribution function. Separating $\bar{\alpha} p_{jst}$ in the mean utility from $-\sigma\nu_i p_{jst}$ in the individual-specific component ensures that $\bar{\alpha}$ is the mean of the individual price coefficients, so that the estimated $\bar{\alpha}$ is directly comparable with the price coefficients from the logit specifications.

The implied own- and cross-price derivatives are

$$\frac{\partial s_{jst}}{\partial p_{kst}} = \begin{cases} \displaystyle\int \alpha_i \, s_{ijst} \left(1 - s_{ijst}\right) d\Phi(\nu), & j = k, \\ \displaystyle -\int \alpha_i \, s_{ijst} \, s_{ikst} \, d\Phi(\nu), & j \neq k, \end{cases} \tag{5}$$

where s_{ijst} is consumer i's choice probability. These derivatives depend on the joint distribution of choice probabilities and price sensitivity across consumers, not on market shares alone, and they are the objects that the supply-side analysis in Section 5.3 uses to recover markups and marginal costs.

5.3 Supply

Firms set prices in a static multiproduct Bertrand game. Firm f chooses prices for its products $\mathcal{J}_{f,st}$ to maximize

$$\pi_{f,st}(p_t) = M_{st} \sum_{j \in \mathcal{J}_{f,st}} (p_{jst} - mc_{jst}) s_{jst}(p_t). \tag{6}$$

Letting Ω_{st} be the ownership matrix with $\Omega_{jk,st} = 1$ when j and k are priced by the same firm and $\Delta_{st} = \partial s_{st}/\partial p_{st}$ the share Jacobian, the first-order conditions are

$$s_{st} + (\Omega_{st} \circ \Delta'_{st})\,(p_{st} - mc_{st}) = 0, \tag{7}$$

where $\circ$ is the Hadamard product. The Jacobian follows from the demand system: for consumer i, $\partial s_{ijst}/\partial p_{kst} = \alpha_i s_{ijst}(\mathbf{1}\{j = k\} - s_{ikst})$, integrated over the distribution of α_i. Since α_i is common across products for a given consumer, Δ_{st} is symmetric here.

I impose full within-firm internalization and do not estimate conduct parameters. Markups and marginal costs follow from inverting (7):

$$\mu_{st} \equiv p_{st} - mc_{st} = -\left(\Omega_{st} \circ \Delta'_{st}\right)^{-1} s_{st}, \qquad mc_{st} = p_{st} - \mu_{st}. \tag{8}$$

Lerner indices are $L_{jst} = \mu_{jst}/p_{jst}$ and variable profits are $\pi_{jst} = \mu_{jst} q_{jst}$.

Ownership is period-specific. Before the transaction, Imperial's brands and the United Sugars entity are separate; afterwards Imperial's brands are priced by the combined entity, which internalizes substitution across the enlarged portfolio. The inversion imposes no cost change: any movement in recovered marginal costs reflects observed prices, estimated demand derivatives, and the maintained ownership and conduct assumptions.

6 Identification and Estimation

6.1 Demand

The demand parameters are estimated by GMM following the approach of Berry et al. (1995). The parameter vector separates into a linear component, comprising the mean price coefficient, the coefficient on the organic seal rate, and the fixed effects, and a nonlinear component consisting of the single heterogeneity parameter σ. For a candidate value of σ, the BLP contraction mapping recovers the vector of mean utilities δ_{jst} that equates the

market shares predicted by equation (4) with the shares observed in the data. Conditional on those mean utilities, the linear parameters are obtained from the instrumental variables moment conditions, and the implied structural error ξ_{jst} follows as the residual. The GMM objective is then evaluated as a quadratic form in the sample moments, and σ is chosen to minimize it.

Identification of the price coefficient requires variation in prices that is unrelated to the unobserved demand shock ξ_{jst}. Retail prices fail this requirement because products with favorable unobserved characteristics tend to command both higher prices and larger shares, which biases the estimated price coefficient toward zero. I therefore use four excluded instruments, chosen to capture cost-side variation and variation in the competitive environment that a product faces. The first is a leave-one-out price, defined as the average price of the same brand and sugar type in other states in the same quarter. Because refined sugar is produced at a small number of facilities and distributed nationally, cost shocks are common across states, so prices elsewhere are informative about the cost of supplying the focal market while excluding that market's own price. The identifying assumption is that demand shocks are not common across states within a brand and quarter, which the quarter and sugar-type by quarter-of-year fixed effects are included. Thus, national movements in demand, whether from input-cost inflation, the shift toward at-home consumption in 2020, or the holiday baking season, are absorbed before the instrument is used.

The second and third instruments are the world raw sugar futures price interacted with sugar-type indicators. Raw sugar is the principal input into refined sugar, so movements in its price shift the marginal cost of supplying retail products while being determined in world markets rather than by demand conditions in any individual U.S. state. The futures price itself varies only over time and is therefore absorbed by the quarter fixed effects, so identification comes from differences across sugar types in the extent to which cost shocks are reflected in retail prices. Those differences are substantial. In the first stage, a one-cent increase in the futures price raises granulated prices by 1.25 cents more than brown sugar

prices, with a standard error of 0.52. Granulated sugar is the product form closest to refined sugar itself, whereas brown and powdered sugar involve additional processing and packaging whose costs move independently of the raw sugar price, so their retail prices respond less strongly to a given cost shock.

The fourth instrument is a differentiation instrument in the spirit of Gandhi and Houde (2019), constructed as $\sum_{k \neq j} (\hat{p}_k - \hat{p}_j)^2$, where $\hat{p}$ denotes the price predicted from the exogenous variables and the excluded cost instruments. This measure summarizes how isolated a product is in price space. It is large when the other products in a market are priced far from product j and small when j competes in a crowded price neighborhood. It is included because the heterogeneity parameter σ is identified by variation in the competitive environment rather than by variation in prices alone. A product with many similarly priced rivals faces more elastic demand than the logit would predict, and the magnitude of that difference is informative about the dispersion of price sensitivity. Because the measure is constructed from predicted rather than observed prices, it depends only on exogenous variation. Together with the exogenous product characteristic and the full set of fixed effects, these four instruments provide the variation used to identify the demand parameters.

6.2 Supply

The estimated demand system produces, for each market, the vector of predicted market shares and the full matrix of own- and cross-price derivatives given in equation (5). These derivatives are the link between the demand estimates and the supply side. They capture not only how a product's own demand responds to its price, but also how that price change diverts demand across the other products in the market, including products sold by the same brand in other sugar types and products sold by competing brands. Because the derivatives depend on the joint distribution of choice probabilities and price sensitivity across consumers rather than on market shares alone, diversion is concentrated among products that appeal to similar consumers.

On the supply side I assume that firms compete in prices according to a static multi-product Bertrand–Nash equilibrium. Each firm chooses the prices of the products it owns to maximize variable profits, taking rivals' prices as given and internalizing the effect of each price on the demand for every product in its own portfolio. I impose full within-firm internalization rather than estimating separate conduct parameters, so the ownership matrix alone determines which cross-product demand effects enter a firm's first-order conditions.

Given the estimated demand derivatives and the ownership structure prevailing in each period, markups are recovered by inverting the firms' first-order conditions. Writing Δ_{st} for the matrix of price derivatives and Ω_{st} for the ownership matrix, the vector of markups satisfies

$$p_{st} - mc_{st} = -\left(\Omega_{st} \circ \Delta_{st}'\right)^{-1} s_{st},$$

and marginal costs are recovered as the residual between observed prices and the implied markups,

$$mc_{st} = p_{st} + \left(\Omega_{st} \circ \Delta_{st}'\right)^{-1} s_{st}.$$

This inversion requires no parametric assumption about the cost function. Marginal costs are recovered product by product from observed prices, estimated demand derivatives, and the maintained conduct assumption, and are therefore free to vary across brands, sugar types, states, and quarters. The ownership matrix changes between periods, which alters the mapping from prices and demand derivatives into implied markups, but any resulting movement in recovered costs is an output of the procedure rather than an assumption.

The recovered marginal costs and markups are used to construct the Lerner index $L_{jst} = (p_{jst} - mc_{jst})/p_{jst}$ and variable profits $\pi_{jst} = (p_{jst} - mc_{jst})q_{jst}$, which provide the basis for evaluating producer margins across brands and over time. Together with the consumer surplus implied by the demand system, they also provide the inputs to the counterfactual analysis, in which equilibrium prices are recomputed under alternative ownership and cost configurations while the estimated demand system is held fixed.

The structural framework therefore complements the reduced-form evidence in two respects. First, it separates observed prices into a marginal cost and a price–cost margin, so that a period in which prices do not rise can be distinguished from one in which margins do not rise. Second, it supplies the demand derivatives and ownership structure needed to simulate counterfactual equilibria, which allows the effect of the ownership change to be isolated from the other determinants of retail prices that moved over the sample period.

7 Results

7.1 Demand

Table 4 reports four demand specifications, all estimated on the same 88,270 product–market observations with the same fixed effects and, where instrumented, the same instruments. Column (1) is a plain logit estimated without instruments, column (2) instruments for price, column (3) adds a nesting structure in which products are grouped by sugar type, and column (4) is the random-coefficients specification used throughout the remainder of the analysis. The first three columns are reported to place the preferred estimates in context. Instrumenting increases the absolute price coefficient by 53 percent, from -0.058 to -0.089, in the direction the standard argument predicts. Products with favorable unobserved characteristics command both higher prices and larger shares, so the uninstrumented estimate understates price sensitivity. Both logit columns nonetheless imply median own-price elasticities above -1 in absolute value, -0.39 and -0.59, as does the nested logit at -0.57. Demand that inelastic is inconsistent with profit maximizing pricing, since a firm facing it could raise its price and increase revenue while selling less.

Table 4: Demand Estimates

	(1)	(2)	(3)	(4)
	Logit	IV Logit	IV Nested	IV RC
Price, $\bar{\alpha}$ (cents/oz)	-0.058***	-0.089***	-0.025***	-0.284***
	(0.004)	(0.007)	(0.003)	(0.039)
USDA organic seal rate	-0.113	-0.039	-0.036***	-0.053
	(0.088)	(0.095)	(0.011)	(0.086)
Nesting parameter ρ			0.723***	
			(0.027)	
Heterogeneity σ				0.132***
				(0.017)
Median own-price elasticity	-0.39	-0.59	-0.57	-1.44
Implied value of organic seal (cents/oz)	-1.95	-0.44	-1.44	-0.19
Brand, type, state, quarter FE	Yes	Yes	Yes	Yes
Type × season FE	Yes	Yes	Yes	Yes
Observations	88,270	88,270	88,270	88,270

Notes: Standard errors clustered by brand and quarter. Columns (2)–(4) instrument for price with the leave-one-out price of the same brand and sugar type in other states, the world raw sugar futures price interacted with sugar type, and a predicted-price differentiation instrument. The implied value of the organic seal is the organic coefficient divided by the absolute value of the mean price coefficient. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Column (4) is the preferred specification. The mean price coefficient is -0.284 with a standard error of 0.039, roughly three times the instrumented logit estimate, and the heterogeneity parameter is 0.132 with a standard error of 0.017. Both are precisely estimated, so the data reject a common price coefficient across consumers. The implied median own-price elasticity is -1.44: a brand raising the prices of its products by 10 percent loses about 14 percent of its volume, against 6 percent under the instrumented logit. Elasticities are

dispersed across products, with an interquartile range of -1.74 to -1.07 and a fifth to ninety-fifth percentile range of -2.03 to -0.46. Weighting by volume gives -1.01, reflecting the fact that the largest products face less elastic demand than the typical one.

The reason heterogeneity matters so much is the substitution restriction that logit imposes. Under logit, a price increase diverts demand to rivals in proportion to their market shares, whether or not those rivals are close substitutes, and a product's own-price elasticity is bounded by $\bar{\alpha} p_j (1 - s_j)$. With a price coefficient of -0.089 and an average price of 4.56 cents per ounce, that bound falls short of unity for nearly every product in the sample. The random coefficient breaks the restriction. The consumers who leave a product when its price rises are disproportionately the price-sensitive ones, and they switch disproportionately to cheaper alternatives. The estimated dispersion is substantial. The price coefficient of a consumer at the tenth percentile of the distribution is -0.447, nearly four times that of a consumer at the ninetieth percentile, -0.120. Individual products therefore face considerably more elastic demand than logit allows, while the category as a whole remains inelastic, which is what one expects of a staple with few close substitutes outside it.

The coefficient on the organic seal rate is -0.053 and imprecisely estimated. Dividing it by the absolute value of the mean price coefficient converts it into a willingness to pay of -0.19 cents per ounce, about 3 cents per pound, against an average retail price of 4.56 cents per ounce. Conditional on brand, sugar type, state, quarter, and the seasonal controls, organic products therefore carry no measurable utility premium. The estimate is a conditional association rather than the causal effect of the label. Organic status is concentrated in a small number of brands, so the brand fixed effects absorb most of its variation, and the characteristic is not itself instrumented.

Table 5 reports average own- and cross-price elasticities for the ten largest brands by volume, aggregated to the brand level so that brands offering several sugar types are treated as portfolios. Rows identify the brand whose aggregate share responds and columns the brand whose prices change, so entry (j, k) is the percentage change in brand j's share from

a one-percent proportional increase in the prices of brand k's products.

The diagonal entries are negative throughout and vary substantially across brands, from -1.43 for Imperial to -0.63 for the largest store brand. The national brands are systematically more elastic than the leading store brands, averaging -1.24 against -0.94. Two forces produce this. The largest store brands hold much larger shares in their local markets, which mechanically lowers the own-price elasticity through the $(1 - s_j)$ term, and their lower prices mean that a one-percent increase is a smaller absolute change to which consumers respond less. Imperial, which is both smaller and higher priced than the leading private labels, sits at the elastic end of the range.

The off-diagonal entries are positive throughout, confirming that the products are substitutes, and average 0.106. Their pattern is more informative than their average. Diversion runs predominantly toward the largest products rather than being spread evenly: the column for the largest store brand contains the biggest entries in most rows, reaching 0.324 for another store brand and 0.254 for Domino, and the largest single entry in the table is the 0.393 response of one store brand to C&H's prices. For Domino, 76 percent of total diversion to the other nine brands goes to private labels. A price increase by a national brand therefore sends consumers primarily to store brands, which is the substitution pattern the random coefficient is designed to capture and which the logit, by construction, cannot generate except in proportion to shares.

The off-diagonal entries are positive throughout, confirming that the products are substitutes, and average 0.106. Their pattern is more informative than their average. Diversion runs predominantly toward the largest products rather than being spread evenly: the column for the largest store brand contains the biggest entries in most rows, reaching 0.324 for another store brand and 0.254 for Domino, and the largest single entry in the table is the 0.393 response of one store brand to C&H's prices. For Domino, 76 percent of total diversion to the other nine brands goes to private labels. A price increase by a national brand therefore sends consumers primarily to store brands, which is the substitution pattern

the random coefficient is designed to capture and which the logit, by construction, cannot generate except in proportion to shares.

Table 5: Average Own- and Cross-Price Elasticities, Ten Largest Brands

	DOMINO	MASKED 130	C&H	MASKED 79	MASKED 5851	MASKED 151	MASKED 32	MASKED 6901	ZULKA	IMPERIAL
DOMINO	**-1.154**	0.254	0.190	0.151	0.035	0.123	0.099	0.020	0.011	0.017
MASKED 130	0.222	**-0.634**	0.036	0.163	0.051	0.160	0.022	0.018	0.010	0.035
C&H	0.125	0.254	**-1.186**	0.129	0.035	0.119	0.098	0.021	0.013	0.020
MASKED 79	0.355	0.324	0.000	**-0.873**	0.036	0.048	0.045	0.016	0.007	0.001
MASKED 5851	0.235	0.289	0.191	0.167	**-0.950**	0.141	0.090	0.023	0.012	0.018
MASKED 151	0.222	0.301	0.039	0.007	0.023	**-0.809**	NA	0.018	0.009	0.001
MASKED 32	0.133	0.103	0.393	0.181	0.025	NA	**-1.294**	0.018	0.018	0.007
MASKED 6901	0.223	0.282	0.210	0.163	0.038	0.139	0.110	**-1.070**	0.012	0.017
ZULKA	0.219	0.278	0.208	0.154	0.037	0.136	0.105	0.022	**-1.197**	0.017
IMPERIAL	0.175	0.231	0.196	0.132	0.033	0.114	0.095	0.020	0.012	**-1.434**

Notes: Averages across state–quarter markets, from the random-coefficients specification in column (4) of Table 4. Rows identify the brand whose aggregate share responds; columns identify the brand whose prices change. Diagonal entries are own-price elasticities. NA indicates brand pairs never observed in the same market, which occurs for private labels belonging to different retailers.

7.2 Marginal Costs

After estimating demand, I recover marginal costs by inverting firms' first-order conditions under the maintained multiproduct Bertrand pricing model. For each product–market observation, the recovered marginal cost is the difference between the observed retail price and the markup implied by the estimated demand system and the assumed ownership structure. Marginal costs are therefore recovered separately for each product in each state and quarter, allowing them to vary across brands, geographic markets, and time.

The ownership matrix used in the supply-side inversion reflects the ownership structure prevailing in each period. Before the merger, products associated with U.S. Sugar and Imperial Sugar are assigned to separate firms. Following the acquisition, Imperial's products are assigned to the combined firm, which internalizes the effect of pricing decisions across its expanded product portfolio. This change in ownership affects the markups implied by the pricing first-order conditions and therefore also affects the recovered marginal costs. Importantly, the inversion does not impose a cost reduction following the merger. Any change in recovered marginal costs therefore reflects the combination of observed prices, estimated demand derivatives, and the maintained ownership and conduct assumptions.

	Price (cents/oz)	Average marginal cost (cents/oz)	Median marginal cost (cents/oz)
All brands	6.47	1.59	2.42
Non–store brands	7.27	1.79	3.69
Store brands	5.59	1.37	1.80

Table 6: Marginal Costs by Brand Type

As reported in Table 6, the volume-weighted average marginal cost across all observations is 1.59 cents per ounce, while the unweighted median is 2.42 cents per ounce. The average gives greater weight to observations with larger quantities sold, whereas the median treats

each product–market observation equally. The lower volume-weighted average therefore indicates that a substantial share of observed sales volume is associated with products having relatively low recovered marginal costs.

A clear difference emerges between store and non-store brands. The median recovered marginal cost of non-store brands is 3.69 cents per ounce, compared with 1.80 cents per ounce for store brands. The ordering is the same on a volume-weighted basis, at 1.79 against 1.37 cents per ounce, although the gap is narrower because the largest store-brand products account for a disproportionate share of volume. These differences may reflect procurement arrangements, distribution, packaging, or product positioning; the structural estimates do not separately identify these mechanisms, so the lower recovered costs of store brands should not by themselves be read as evidence of greater productive efficiency.

Recovered costs are more similar across sugar types than across brand types. The median is 2.57 cents per ounce for granulated sugar, 2.59 for brown, and 2.20 for powdered, with median Lerner indices between 0.66 and 0.68. Because granulated sugar is both the cheapest per ounce at retail and the largest category by volume, its volume-weighted average marginal cost, 1.34 cents per ounce, is the lowest of the three.

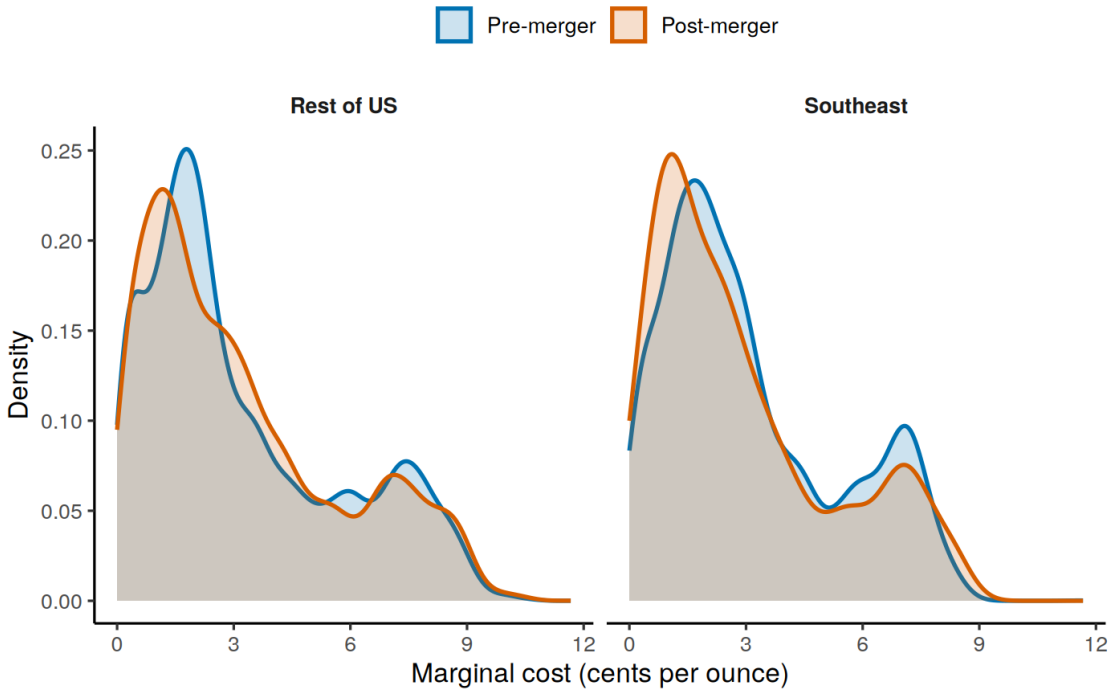


Figure 6: Distribution of Marginal Costs Before and After the Merger

Figure 6 compares the distributions of recovered marginal costs before and after the merger. The two distributions overlap substantially over the range where most observations are concentrated. The median recovered marginal cost is 2.39 cents per ounce before the merger and 2.49 cents after, while the volume-weighted average declines slightly from 1.62 to 1.53 cents per ounce. The unconditional distributions therefore do not indicate a simple, uniform shift in marginal costs following the merger.

This comparison is descriptive and does not identify a causal effect of the merger on costs. Product and market composition change over time through entry, exit, and shifts in geographic availability, and the 2021–22 period was marked by broad increases in food and input costs. Moreover, because the supply-side inversion incorporates the contemporaneous ownership structure, the change in ownership itself alters the mapping from observed prices and estimated demand derivatives to implied markups and costs. The counterfactual analysis later instead holds the demand environment fixed and varies ownership and costs separately.

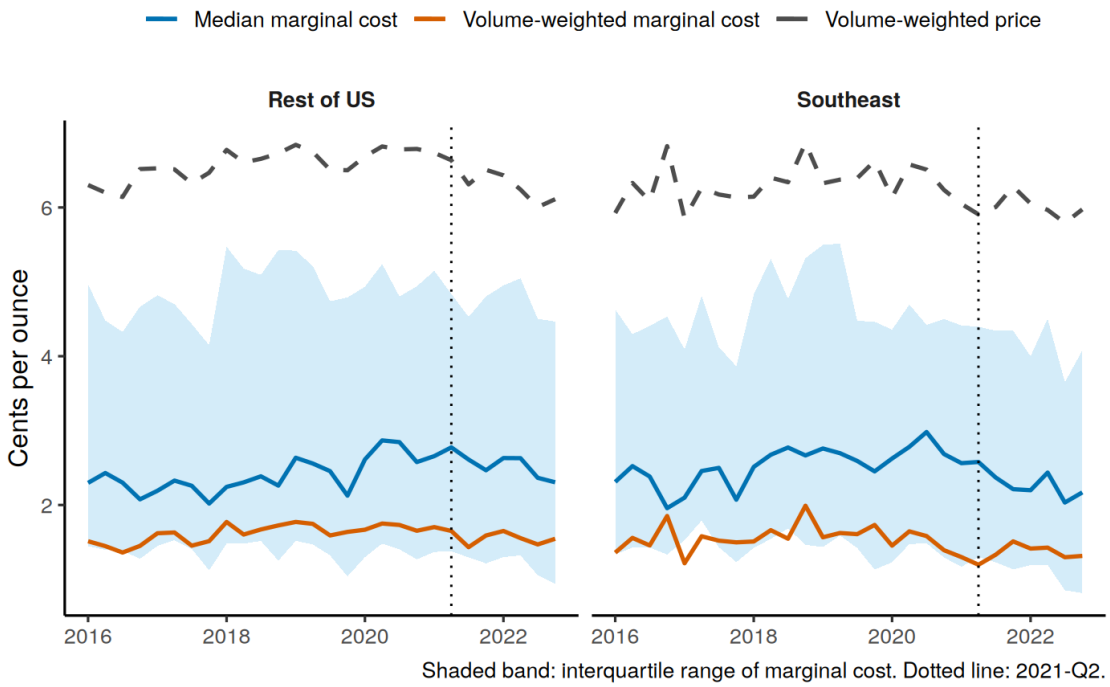


Figure 7: Recovered Marginal Costs and Prices over Time

Figure 7 plots the quarterly series separately for the Southeast and the rest of the country. Recovered marginal costs track retail prices upward through 2021–22, which is what would be expected if rising input costs were passed through, and neither region displays a discrete break at the merger quarter. In the Southeast, the median recovered marginal cost declines from 2.53 to 2.26 cents per ounce across the two periods while the median Lerner index rises slightly, from 0.667 to 0.679.

Taken together, the results reveal three main patterns. First, recovered marginal costs vary substantially across products and markets, with a volume-weighted average of 1.59 cents per ounce and a median of 2.42. Second, store brands have materially lower recovered marginal costs than non-store brands under both measures. Third, the pre- and post-merger distributions overlap substantially and do not indicate a uniform change in marginal costs following the merger. These recovered costs provide the basis for the markups and Lerner indices reported next and for the counterfactual decomposition that separates ownership from cost effects.

7.3 Markups and the Lerner Index

Using the recovered marginal costs, I calculate a markup and Lerner index for each product–market observation. The markup is the difference between price and marginal cost, measured in cents per ounce, while the Lerner index normalizes this difference by the product's price. The two measures capture related but distinct dimensions of pricing incentives. The markup measures the absolute price–cost margin, whereas the Lerner index measures that margin relative to price. Table 7 reports average markups, average Lerner indices, and the correlation between markups and prices for all brands and separately for non-store and store brands.

	Average Markup	Median Markup	Average Lerner	Corr(Markup, Price)
All brands	6.25	4.44	68.4%	0.931
Non–store brands	7.87	5.78	66.8%	0.930
Store brands	4.44	4.07	70.3%	0.899

Table 7: Markup and Lerner Index by Brand Type

Across all brands, the average recovered markup is 6.25 cents per ounce and the average Lerner index is 68.4 percent, so the estimated margin accounts for roughly two-thirds of the retail price. The average exceeds the median of 4.44 cents because the markup distribution has a long right tail.

The average markup differs substantially across brand types. Non-store brands have an average markup of 7.87 cents per ounce against 4.44 for store brands, a difference of 3.43 cents, or 77 percent relative to the store-brand average. Non-store brands therefore earn considerably larger absolute price–cost margins, consistent with their higher prices and higher recovered marginal costs.

The ranking reverses when margins are expressed relative to price. The average Lerner index is 66.8 percent for non-store brands and 70.3 percent for store brands. Although store

brands earn smaller absolute markups, those markups represent a slightly larger fraction of their lower retail prices. Absolute markups and proportional margins therefore give different pictures of pricing across the two categories, and both are reported below.

The final column shows that markups and prices are strongly positively correlated: 0.931 across all brands, 0.930 among non-store brands, and 0.899 among store brands. Because the markup is the difference between price and a recovered cost with smaller variance, this correlation is largely mechanical and indicates that variation in prices, rather than in recovered costs, drives most of the variation in absolute margins.

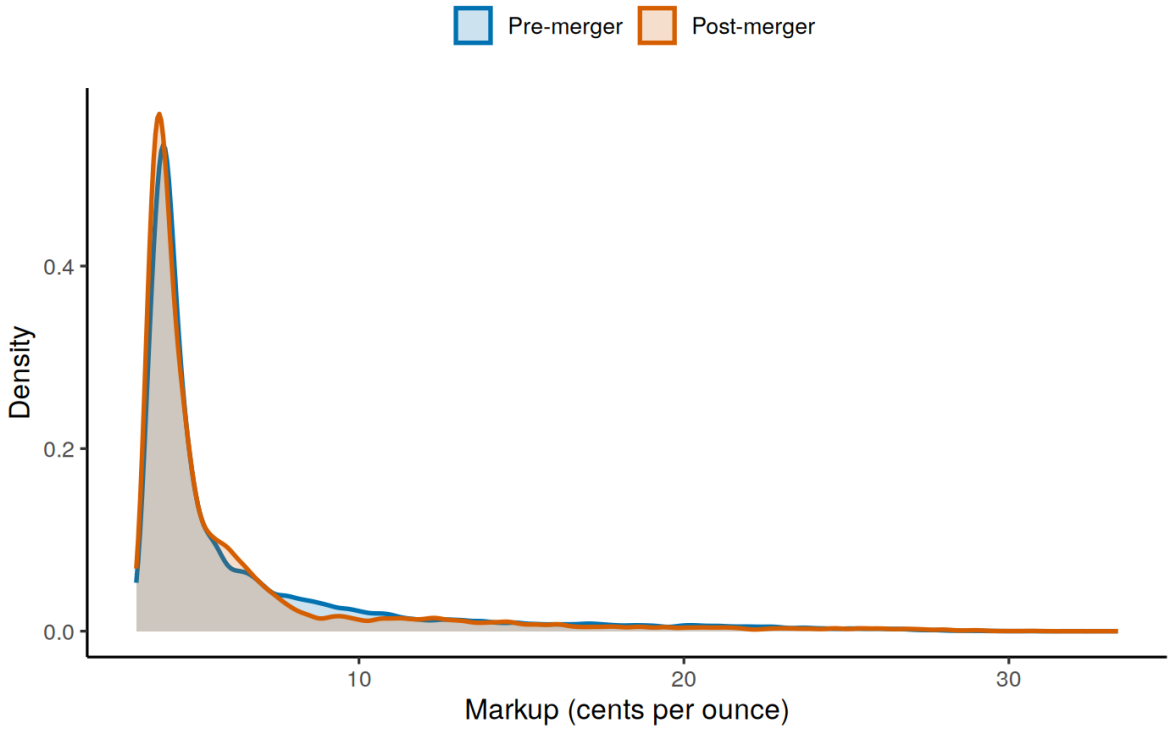


Figure 8: Distribution of Markups Before and After the Merger

Figure 8 compares the distributions of product–market markups before and after the merger. Both are concentrated between roughly 2 and 8 cents per ounce with long right tails, and they overlap over most of their support. Median markups are 4.47 cents per ounce before the merger and 4.35 cents after. The figure does not indicate a discrete increase in markups following the transaction.

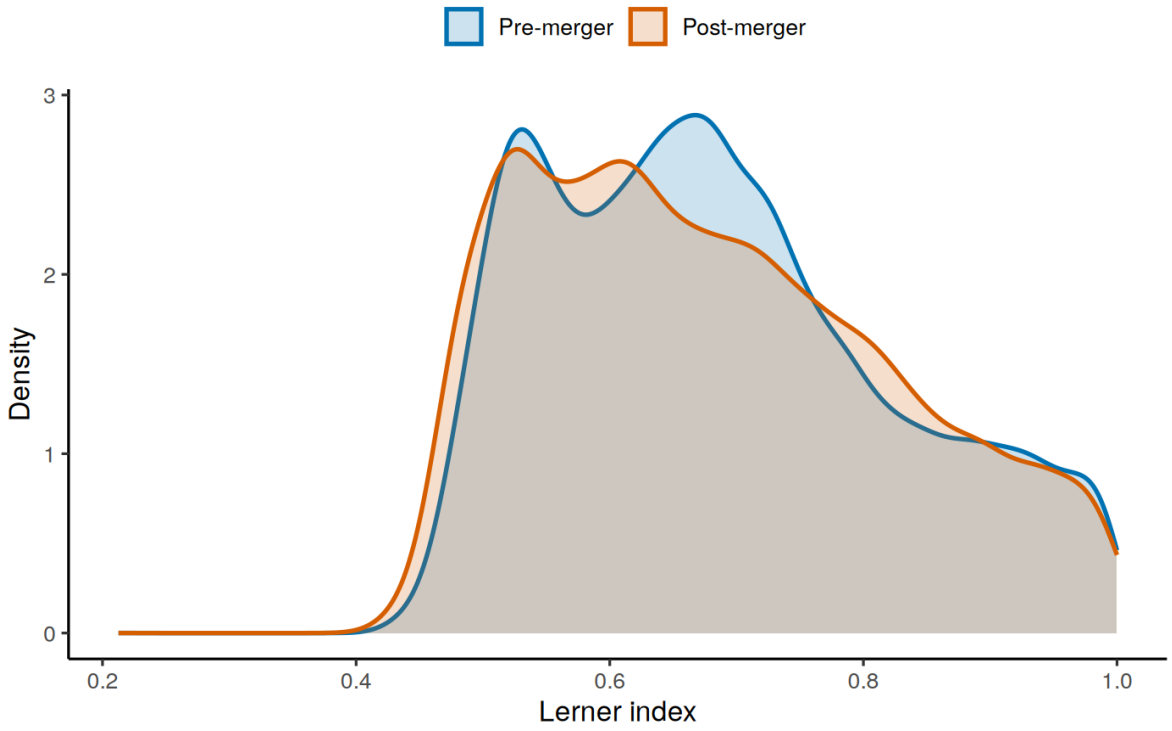


Figure 9: Distribution of Lerner Indices Before and After the Merger

Figure 9 presents the corresponding distributions of Lerner indices. Because the index divides the markup by price, it is bounded between zero and one in the recovered sample and is directly comparable across products at different price levels. The two distributions again overlap substantially, with median values of 0.670 before and 0.662 after the merger.

The contrast between the two figures is informative. Because the Lerner index divides the markup by price, the two measures need not move together. A product can experience a change in its absolute markup without a proportional change in its Lerner index if its price moves at the same time. Taken together, neither measure indicates a broad increase in price–cost margins following the merger.

7.4 Marginal Costs, Markups, and Lerner Indices for Acquired Brands

This subsection traces the two brands acquired in the transaction, Imperial and Dixie

Crystals, and compares them with the rest of the category. Each series is a volume-weighted average across the markets in which the brand is observed in a given quarter, computed separately for the Southeast and the rest of the country and reported alongside three comparison groups, the brands of American Sugar Refining, other national brands, and private label. The vertical dashed line marks 2021Q2.

Because the set of markets underlying each point changes over time, I first document the brands' retail footprint in Figure 10. Both brands expand it over the sample. Imperial is observed in 25.1 markets per quarter before the announcement and 34.4 after; Dixie Crystals in 15.8 and 28.1 respectively. The number of products each brand offers is unchanged at six, so the expansion reflects wider geographic availability rather than a broader assortment, and there is no sign of the product rationalization that a merged firm might undertake. The brands' regional presence differs markedly. Dixie Crystals is observed in 7.0 Southeast markets per quarter after the announcement against 28.1 elsewhere, while Imperial is observed in 7.0 and 34.4 respectively; measured by volume, the Imperial brand is only marginally present in the region containing the refinery it is named for.

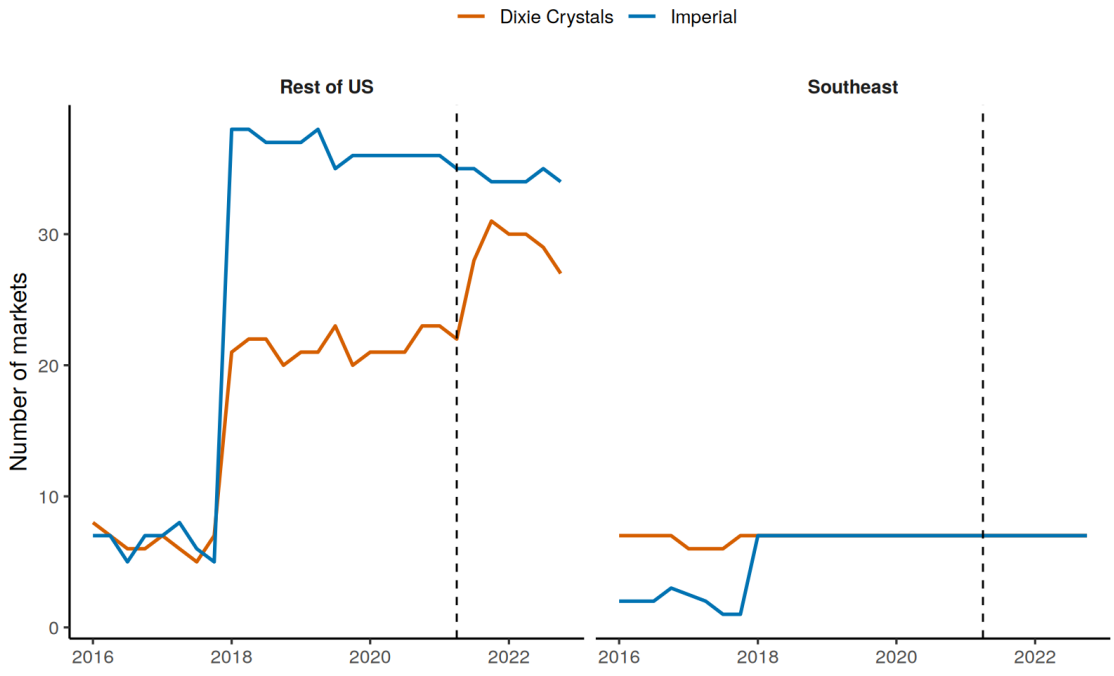


Figure 10: Markets in Which the Acquired Brands Are Observed

Figure 11 shows retail prices, the one series observed rather than model-implied. Both acquired brands are priced below other national brands throughout and close to private label. Prices decline for both, by 1.2 percent for Imperial and 7.4 percent for Dixie Crystals between the two periods, and the declines are not specific to them: other national brands fall 11.6 percent and private label 5.0 percent over the same window, while ASR's brands are essentially flat. Neither acquired brand exhibits a discrete change at 2021Q2, and within the Southeast both decline further than they do elsewhere, with Dixie Crystals falling from 4.86 to 4.57 cents per ounce and Imperial from 7.08 to 6.47.

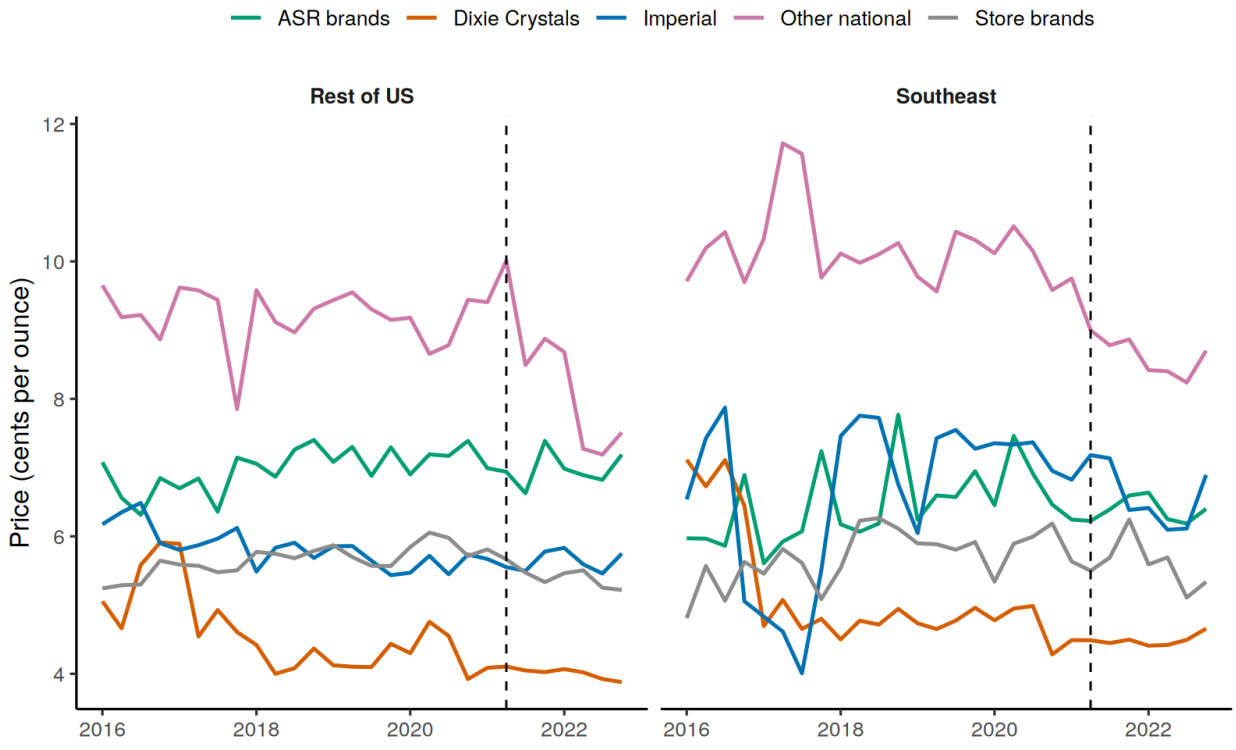


Figure 11: Prices for the Acquired Brands and Comparison Groups

Figure 12 reports recovered marginal costs. The acquired brands have among the lowest recovered costs in the category, between roughly 0.5 and 2.0 cents per ounce, well below the 3.0 to 3.7 cents of other national brands. Their movements across the two periods run in opposite directions: Imperial's recovered cost rises 23.7 percent while Dixie Crystals' falls 31.9 percent. The comparison groups show that neither is unusual. ASR's recovered costs

rise 27.0 percent over the same period while those of other national brands and private label fall 17.1 and 12.7 percent.

The regional split reveals a pattern the pooled series conceals. Outside the Southeast the acquired brands' recovered costs rise, from 0.79 to 0.85 cents per ounce, whereas within the Southeast they fall, from 0.82 to 0.62. The two regions therefore move in opposite directions for the same brands, which is difficult to reconcile with a common change in the cost of supplying them and points instead to changes in the mix of products and markets in which they are sold.

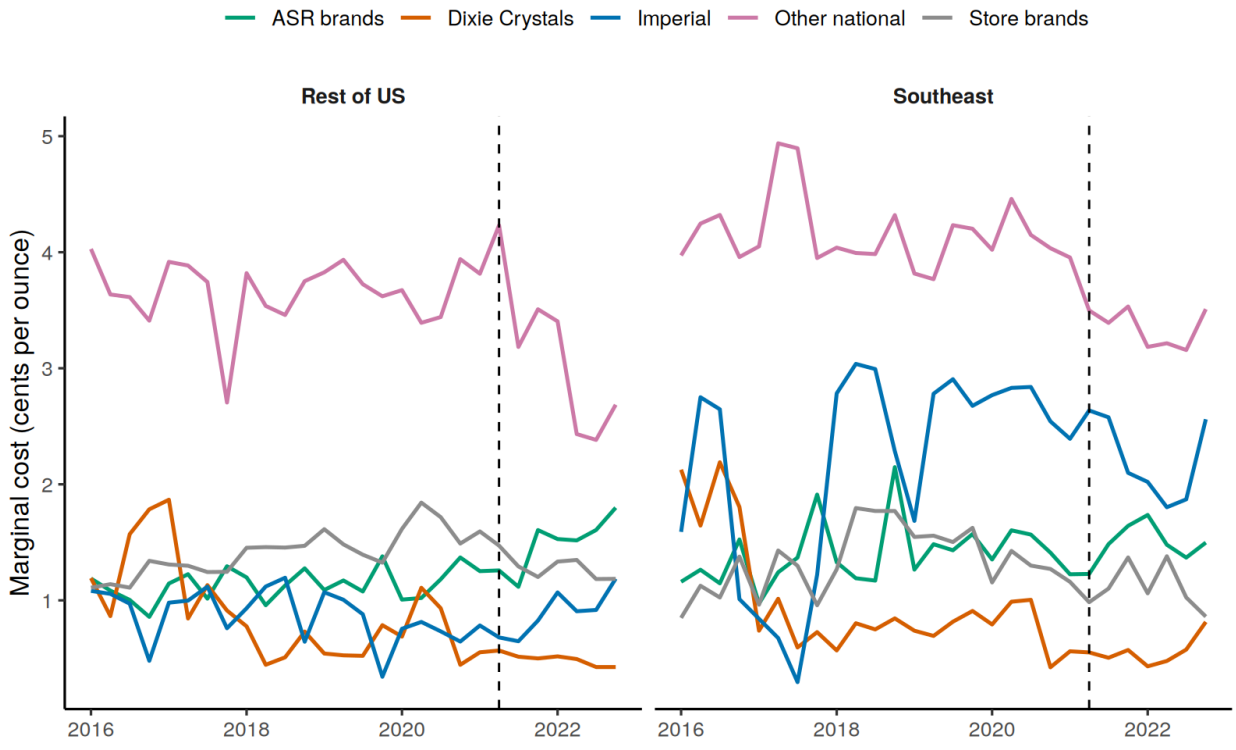


Figure 12: Recovered Marginal Costs for the Acquired Brands and Comparison Groups

Figures 13 and 14 report absolute and proportional margins, and Table 8 summarizes the changes across all five groups. Markups decline for the acquired brands, by 5.3 percent for Imperial and 2.5 percent for Dixie Crystals. They also decline for every comparison group, and by more for other national brands, at 7.9 percent, and ASR's brands, at 6.0 percent, than for either acquired brand. Margins in this category were compressing over the post-

announcement period generally, and the acquired brands compressed by slightly less than their closest comparators rather than expanding.

	Price		Marginal cost		Markup	
	Pre	Post	Pre	Post	Pre	Post
Imperial	5.78	5.71	0.81	1.00	4.97	4.70
Dixie Crystals	4.77	4.42	0.80	0.55	3.97	3.87
ASR brands	6.93	6.91	1.19	1.51	5.75	5.40
Other national	9.24	8.17	3.68	3.05	5.56	5.12
Store brands	5.68	5.40	1.43	1.25	4.25	4.15

Table 8: Prices, Marginal Costs, and Markups by Brand Group (cents per ounce, volume weighted)

The Lerner index tells a slightly different story for Dixie Crystals, rising from 0.721 to 0.761 while its absolute markup falls. The two measures diverge because its price and recovered cost fall together: a markup of 3.87 cents represents a larger share of a price of 4.42 cents than a markup of 3.91 cents represents of 4.77. For Imperial the two measures move together, with the Lerner index rising only slightly, from 0.673 to 0.695. In neither case does the proportional margin rise for a reason that would indicate increased pricing power, and in both cases the changes are of the same magnitude as those observed for the comparison groups.

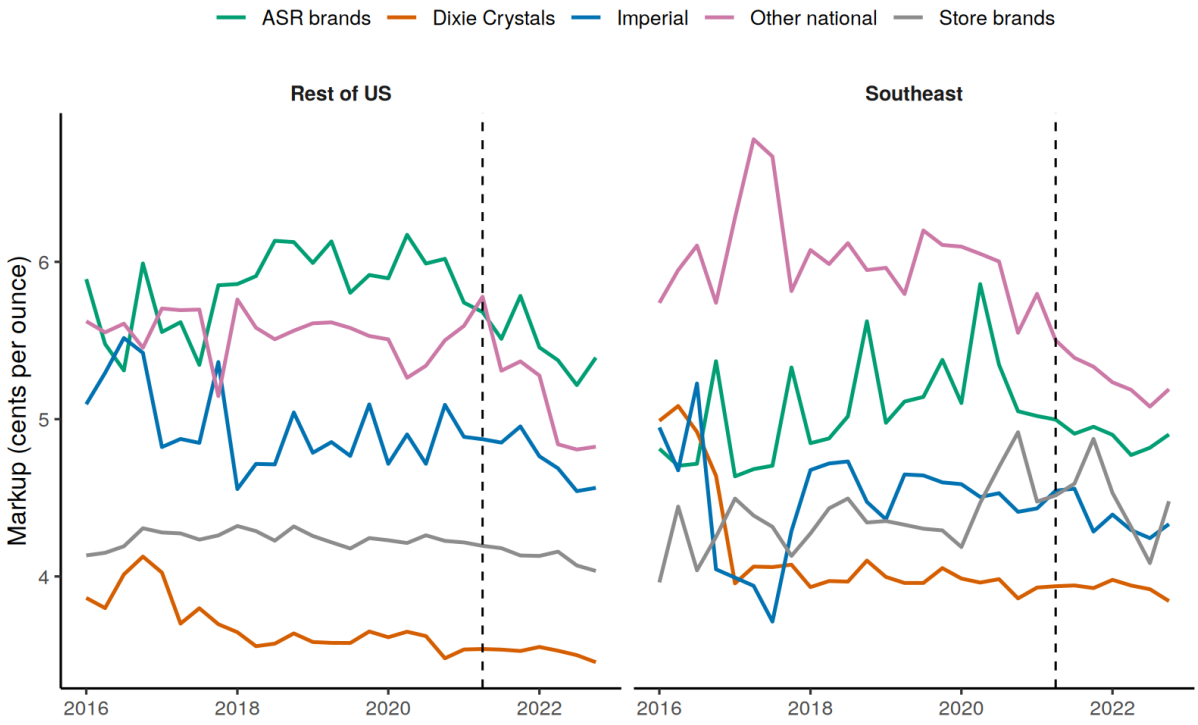


Figure 13: Markups for the Acquired Brands and Comparison Groups

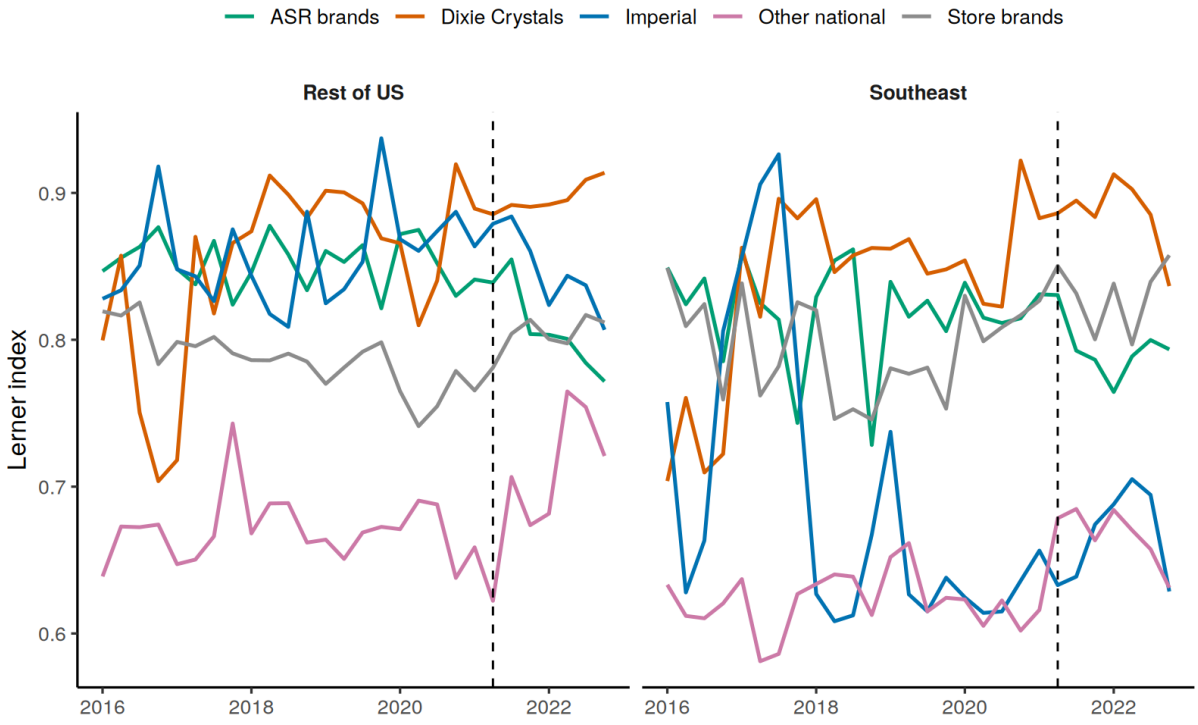


Figure 14: Lerner Indices for the Acquired Brands and Comparison Groups

Overall, the brand-level evidence provides no indication that the transaction expanded the pricing margins of the acquired brands. Their prices decline, their markups decline by slightly less than those of comparable national brands, and their recovered costs move in opposite directions across regions in a way that tracks product and market composition rather than the timing of the transaction. What does change after the announcement is the brands' geographic reach. Both are observed in substantially more markets, which accounts for the increase in their variable profits despite the absence of any increase in per-unit margins.

7.5 Profits and Welfare

I next compare realized market outcomes before and after the merger. Variable profits are recovered price–cost margins multiplied by observed quantities and exclude fixed costs. Consumer surplus is calculated from the estimated demand system using the full set of products available in each market. Total variable welfare is the sum of the two. Because the pre-merger period contains 1,029 state–quarter markets and the post-merger period 343, totals are not comparable across periods, so all measures are reported per market. Table 9 summarizes the averages, and Figures 15–17 show the distributions from which they are computed. Each figure plots kernel densities on a logarithmic scale, since market size varies by several orders of magnitude across state–quarter markets.

	Markets	Variable profit ($m per market)	Consumer surplus ($m per market)	Total welfare ($m per market)
All states, pre-merger	1,029	0.82	9.38	10.19
All states, post-merger	343	1.18	9.02	10.20
Southeast, pre-merger	147	0.70	12.74	13.44
Southeast, post-merger	49	1.10	10.48	11.58

Table 9: Variable Profits, Consumer Surplus, and Total Welfare per Market

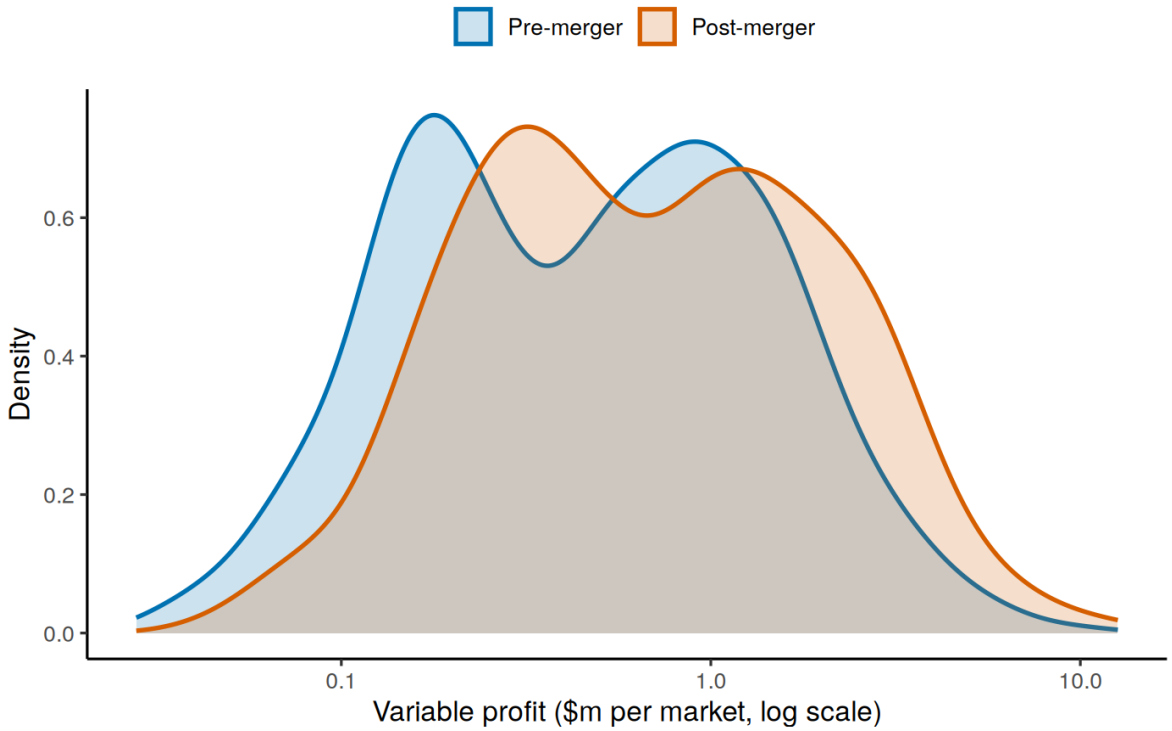


Figure 15: Distribution of Variable Profits Before and After the Merger

Variable profits per market are higher after the merger, averaging $1.18 million against $0.82 million before, an increase of 45 percent. The median rises from $494,000 to $613,000, so the difference is not confined to the upper tail. Figure 15 shows the two distributions. Both are strongly right-skewed, with most markets earning modest variable profits and a small number of large markets accounting for much of the total, and the post-merger distribution is shifted to the right across most of its support rather than differing only in its tail.

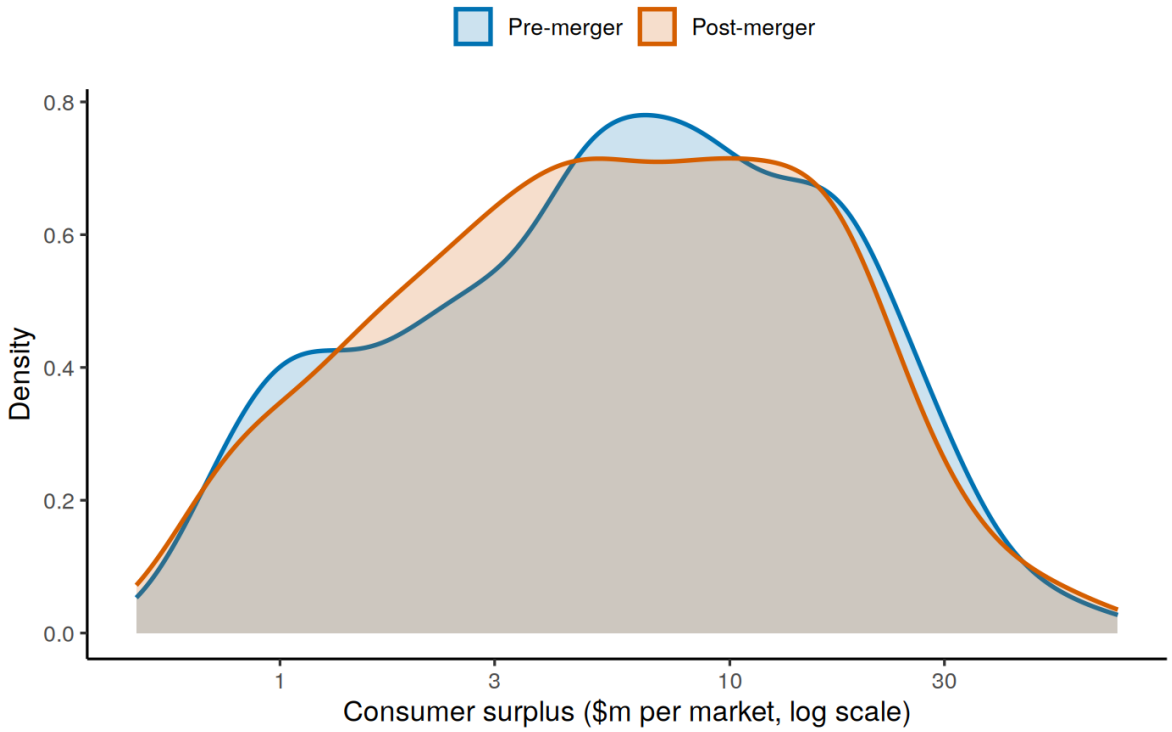


Figure 16: Distribution of Consumer Surplus Before and After the Merger

Consumer surplus moves in the opposite direction, falling from $9.38 million to $9.02 million per market, a decline of 3.8 percent, with the median falling from $6.09 million to $5.65 million. Figure 16 shows that the shift is smaller than the corresponding shift in profits and concentrated in the middle of the distribution, with the two densities nearly coinciding at both ends. The contrast with variable profits is consistent with a modest reallocation of surplus from consumers toward firms over the period.

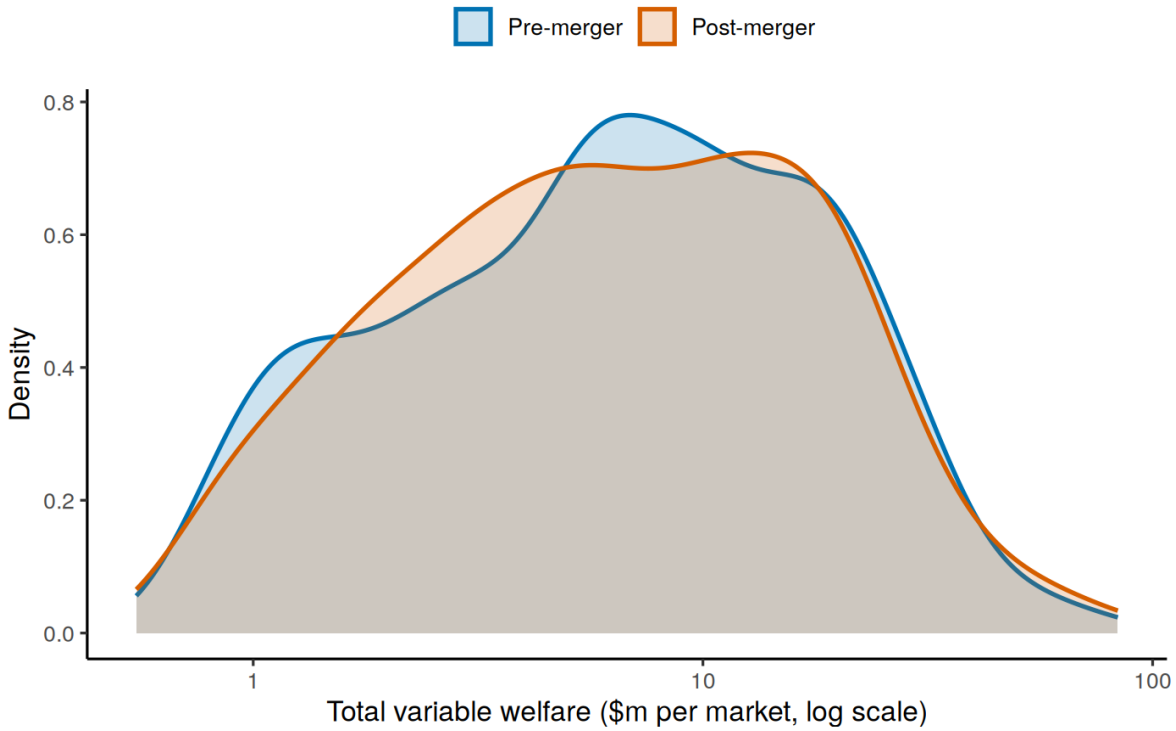


Figure 17: Distribution of Total Welfare Before and After the Merger

Total variable welfare per market is essentially unchanged, at $10.19 million before and $10.20 million after, because the increase in producer surplus and the decline in consumer surplus are of similar magnitude in dollars. Figure 17 shows the two distributions overlapping closely throughout, which is what the offsetting movements in the two components imply. Because consumer surplus is an order of magnitude larger than variable profit in these markets, the welfare distribution closely resembles the consumer-surplus distribution.

These comparisons are descriptive. They combine changes in market size, product composition, demand conditions, input costs, and ownership and they cannot be simply interpreted as effects of the transaction. The counterfactual analysis later holds the demand environment fixed and varies ownership and marginal costs separately for precisely this reason.

7.6 Profit and Welfare in the Southeast

Restricting attention to the region most exposed to the transaction gives a similar picture

with somewhat larger magnitudes. Variable profits per market rise from $0.70 million to $1.10 million, an increase of 56 percent, while consumer surplus falls from $12.74 million to $10.48 million, a decline of 17.8 percent. Total welfare per market falls from $13.44 million to $11.58 million. Both movements are larger than their national counterparts, where profits rise 45 percent and consumer surplus falls 3.8 percent. The Southeast comprises 147 pre-merger and 49 post-merger state–quarter markets, so the distributions below are estimated on far fewer observations than their national counterparts and individual large markets carry correspondingly more weight.

Figure 18 shows the distribution of variable profits in the region. As in the national sample, the post-merger density is shifted to the right across most of its support rather than differing only in the upper tail,and the shift is visibly larger than the national one shown in Figure 15.

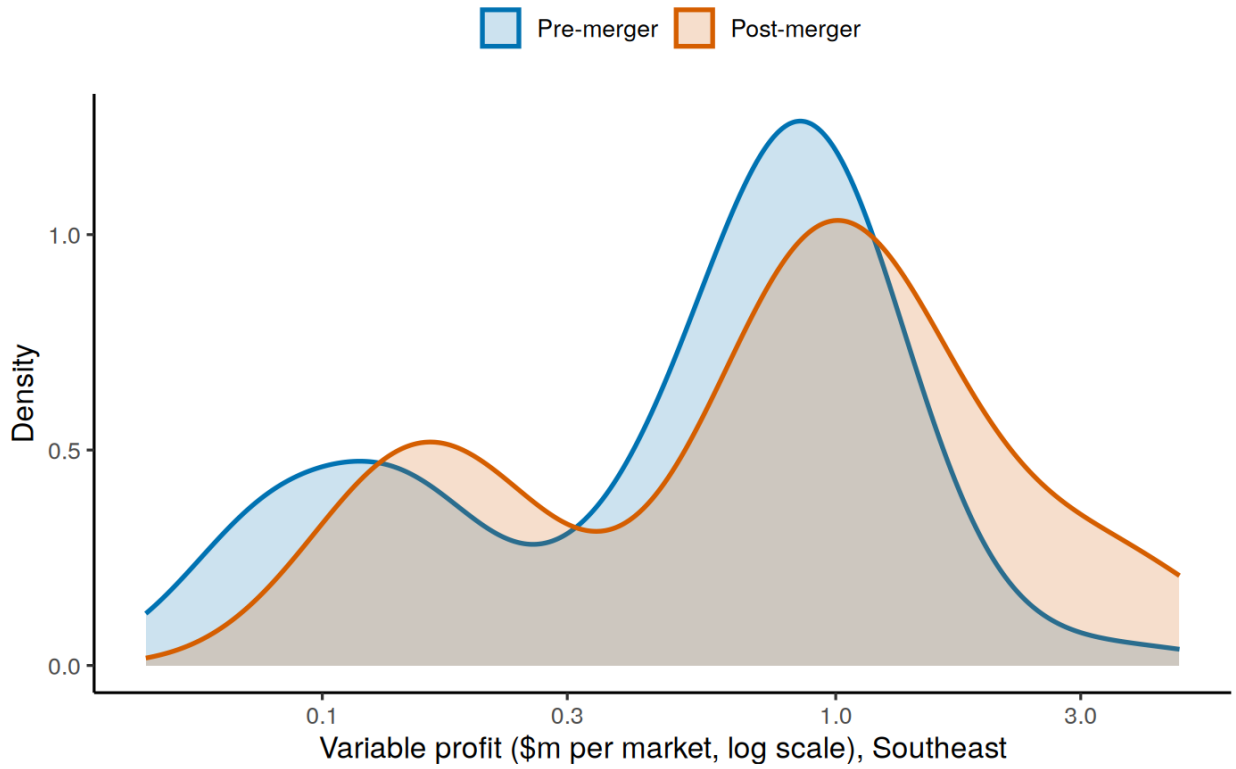


Figure 18: Density of Variable Profits: Pre- vs. Post-Merger in the Southeast

Figure 19 shows total welfare. Unlike the national distributions, which overlap closely,

the post-merger density lies to the left of its pre-merger counterpart over most of the support. Because consumer surplus is roughly an order of magnitude larger than variable profit in these markets, the larger decline in consumer surplus is not offset by the larger increase in profits, and total welfare falls.

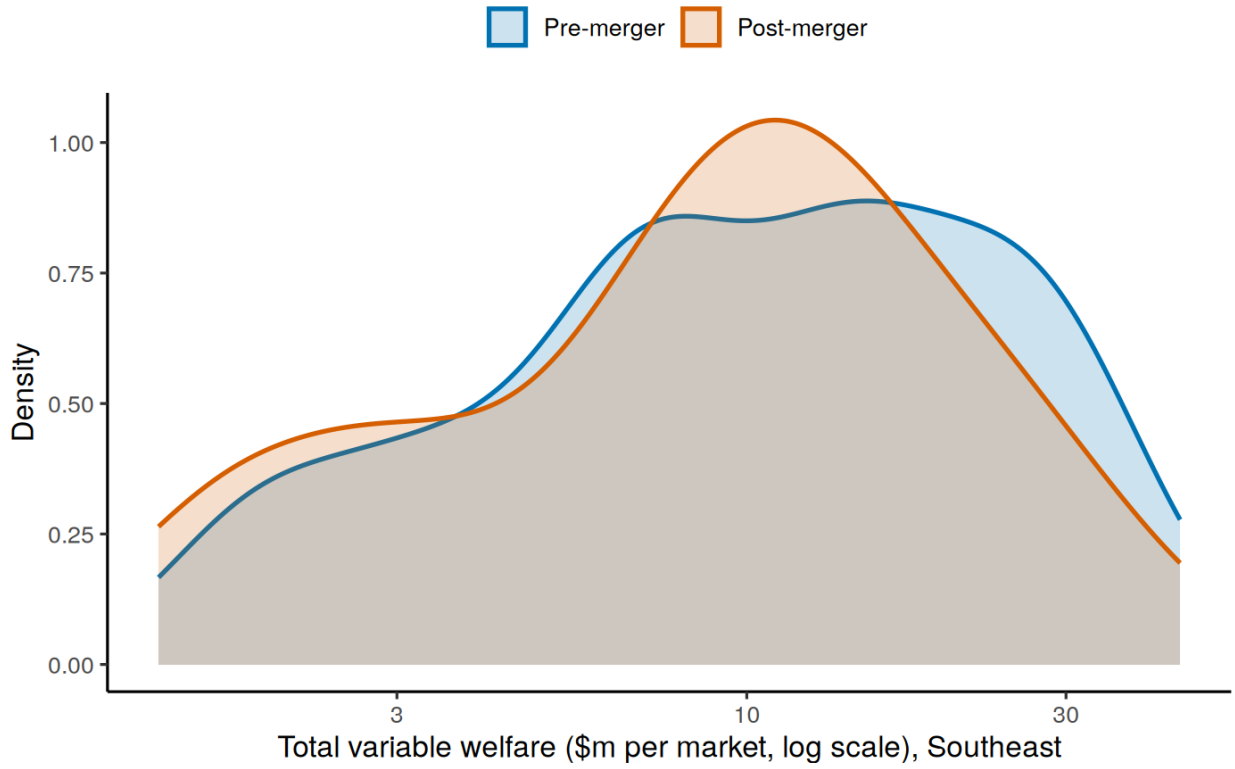


Figure 19: Density of Total Welfare: Pre- vs. Post-Merger in the Southeast

The Southeast therefore shows a larger increase in producer surplus and a larger decline in consumer surplus than the national sample. The pattern is consistent with the region's greater exposure to the transaction, but it is not evidence of it. Variable profit is the product of a margin and a quantity, so profits can rise while relative prices do not, and recovered costs and volumes both move between the two periods. The reduced-form estimates in Section 5 point the same way. Southeast prices rose 2.8 percent less than prices elsewhere rather than more, which is the opposite of what a regional exercise of market power would produce. Separating the contribution of ownership from the other changes in costs, volumes, and product composition requires holding the demand environment fixed, which motivates

the counterfactual decomposition.

8 Policy Analysis

8.1 Setup

To evaluate the economic implications of the U.S. Sugar–Imperial Sugar transaction, I use the estimated demand system and the recovered marginal costs to simulate equilibrium outcomes under alternative ownership and cost configurations. The counterfactual is designed to separate the effect of the change in ownership from the effect of the change in the cost environment, because observed post-merger outcomes reflect both.

I construct four equilibrium scenarios by combining pre- and post-merger ownership structures with pre- and post-merger marginal costs. The first combines pre-merger ownership with pre-merger marginal costs,

$$(\Omega^{pre}, mc^{pre}),$$

and serves as the baseline. The second changes ownership while holding costs fixed,

$$(\Omega^{post}, mc^{pre}).$$

The third changes costs while holding ownership fixed,

$$(\Omega^{pre}, mc^{post}).$$

The fourth combines both,

$$(\Omega^{post}, mc^{post}).$$

All four scenarios are evaluated in the post-merger product–market environment. Estimated demand parameters, product characteristics, mean utilities net of price, and market size are held fixed while ownership and marginal costs vary. Equilibrium prices solve the

firms' Bertrand–Nash first-order conditions in each market and scenario; shares, markups, variable profits, consumer surplus, and total variable welfare follow. The ownership effect is

$$Y(\Omega^{post}, mc^{pre}) - Y(\Omega^{pre}, mc^{pre}),$$

the cost effect is

$$Y(\Omega^{pre}, mc^{post}) - Y(\Omega^{pre}, mc^{pre}),$$

and the total effect compares the post-merger configuration with the baseline,

$$Y(\Omega^{post}, mc^{post}) - Y(\Omega^{pre}, mc^{pre}).$$

For ease of interpretation, I report these effects primarily as percentage changes relative to the pre-ownership, pre-cost benchmark. The percentage measures therefore describe how much the corresponding equilibrium outcome changes when ownership or marginal costs are changed while the remaining components of the post-merger market environment are held fixed.

8.2 Results

Table 10 reports the percentage decomposition of the counterfactual results. The first column reports the effect of changing ownership from the pre-merger to the post-merger ownership structure while holding marginal costs at their pre-merger levels. The second column reports the effect of replacing pre-merger marginal costs with post-merger marginal costs while retaining the pre-merger ownership structure. The final column reports the total change when both ownership and marginal costs are changed

	Ownership Effect (%)	Cost Effect (%)	Total Effect (%)
Price	0.028	-0.444	-0.417
Markup	0.038	-2.135	-2.098
Lerner index	0.010	-1.406	-1.394
Variable profit	0.027	-2.001	-1.975
Consumer surplus	-0.011	-0.052	-0.064
Welfare	-0.003	-0.506	-0.508

Table 10: Decomposition of Counterfactual Merger Effects

Notes: All scenarios are evaluated in the post-merger product–market environment with demand held fixed. Percentage changes are relative to the (Ω^{pre}, mc^{pre}) baseline. Prices, markups, and Lerner indices are weighted by equilibrium quantities; profits, consumer surplus, and welfare are summed across the 343 post-merger markets.

The change in ownership has no economically meaningful effect on equilibrium outcomes. Moving from separate ownership of Imperial's brands and the United Sugars portfolio to joint pricing raises the volume-weighted price by 0.028 percent. Markups rise by 0.038 percent and variable profits by 0.027 percent, while consumer surplus falls by 0.011 percent and total welfare by 0.003 percent. Every effect is smaller than one-twentieth of one percent, and the price effect amounts to roughly three-hundredths of a cent on a five-cent product.

An aggregate this small could in principle conceal larger effects that offset across markets, but it does not. The left panel of Figure 20 shows the distribution of the ownership effect on prices across the 343 post-merger markets. The effect is positive in almost every market, as multiproduct pricing implies, but its magnitude is negligible throughout: the median market sees a price increase of 0.002 percent, the ninety-fifth percentile 0.14 percent, and the largest effect anywhere is 0.56 percent. Only 23 of 343 markets show an increase above one-tenth of one percent. The transaction does raise prices where the merging parties' products meet, and it does so in the direction theory predicts, but by an amount no dataset of this kind could detect.

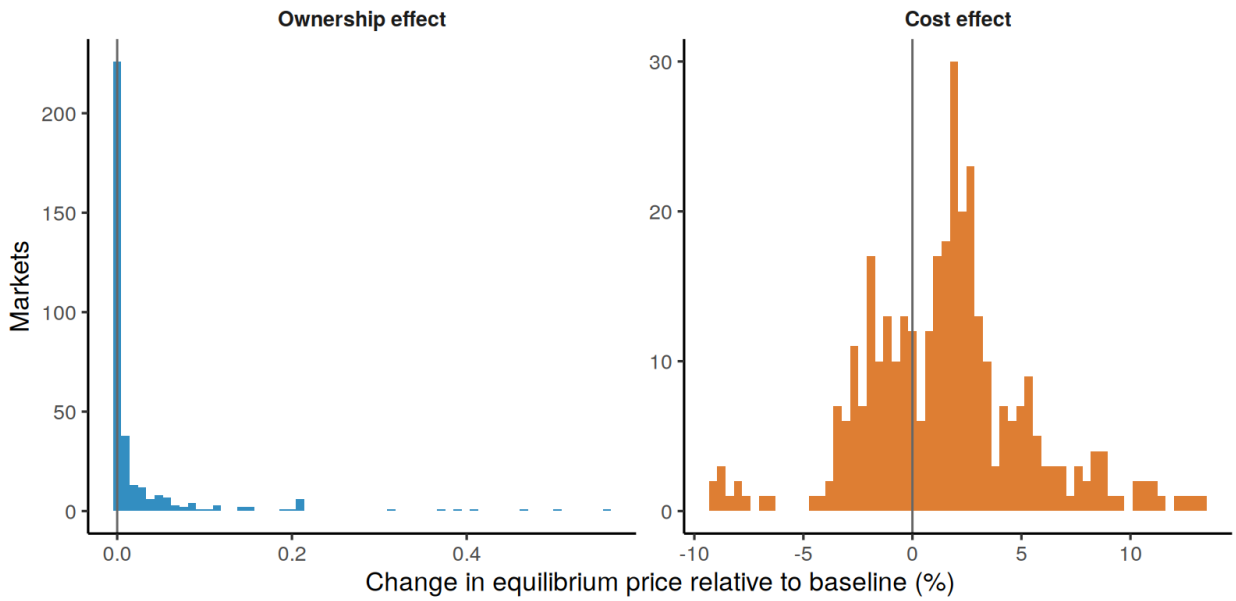


Figure 20: Market-Level Price Effects of the Ownership and Cost Counterfactuals

Two features of the setting explain why. The first is the size of the merging parties' retail presence. Imperial and Dixie Crystals account for 3.3 percent of national volume and U.S. Sugar's retail brands for less than 0.1 percent, so common ownership covers a small share of the category. The second is the pattern of substitution between their products. Under multiproduct Bertrand competition, the unilateral price effect of a merger depends on diversion between the merging firms' products. A price increase is profitable only to the extent that the demand it sheds is recaptured by a product the same firm now owns. The elasticity estimates in Table 5 show that this diversion is small. Imperial's cross-price elasticities with the other large brands are among the lowest in the matrix, and the products brought under common pricing are not close substitutes for one another in the retail market. Little demand is recaptured, so little additional pricing incentive is created.

The cost counterfactual produces larger effects, though they remain modest. Replacing the pre-merger cost environment with the post-merger one lowers equilibrium prices by 0.44 percent, absolute markups by 2.13 percent, and Lerner indices by 1.41 percent, while variable profits fall by 2.00 percent. Consumer surplus is essentially unchanged at -0.05 percent,

and total welfare falls by 0.51 percent.

Unlike the ownership effect, the aggregate here averages over genuinely offsetting movements, shown in the right panel of Figure 20. Prices rise in two-thirds of markets and fall in the remaining third, with the median market seeing an increase of 1.67 percent, the fifth percentile a decline of 3.46 percent, and the ninety-fifth an increase of 8.45 percent. The aggregate is negative because the effect is systematically smaller in larger markets: the correlation between a market's cost effect and its log volume is -0.42, so the unweighted mean of 1.54 percent becomes -0.25 percent once markets are weighted by the quantity they sell. Cost differences between the two periods therefore move prices in both directions depending on the local product mix, which is what one would expect if they reflect changes in composition and input costs rather than a uniform efficiency gain.

Because the ownership effect is negligible, the total effect is almost indistinguishable from the cost effect. Changing both ownership and costs lowers prices by 0.42 percent and markups by 2.10 percent, against 0.44 and 2.13 percent when costs alone change. The results therefore attribute the model-implied change in equilibrium outcomes to marginal costs rather than to ownership. Under the maintained demand and multiproduct Bertrand assumptions, the change in ownership on its own raises equilibrium prices by 0.03 percent.

These results qualify the descriptive comparisons in Section 7, where realized variable profits per market are 45 percent higher after the announcement and consumer surplus 3.8 percent lower. Holding the demand environment fixed and varying ownership alone produces changes three orders of magnitude smaller, so almost none of the observed movement can be attributed to the change in ownership. The descriptive differences reflect the other determinants of retail outcomes discussed there, and cannot be read as evidence of increased market power following the acquisition.

Overall, the counterfactual analysis provides little evidence that the U.S. Sugar–Imperial Sugar transaction generated a substantial unilateral pricing effect through ownership consolidation. Instead, the estimated post-announcement cost environment accounts for nearly

all of the model-implied change in equilibrium outcomes. The result is consistent with a setting in which differences in marginal costs are reflected in prices and margins rather than accompanied by an expansion in price–cost margins.

9 Conclusion

This paper asks whether a merger between sugar refiners raised prices in the retail markets where their brands compete. Using NielsenIQ Retail Scanner data covering 2016 to 2022, I combine reduced-form regional comparisons with a random-coefficients demand model and a multiproduct Bertrand pricing model to recover markups and marginal costs, and I use the estimated model to separate the ownership and cost channels through which the transaction could have affected retail outcomes.

The reduced-form evidence shows no relative price increase in the Southeast, the region containing the refining capacity most directly affected by the acquisition. The estimated differential is -2.8 percent for prices and statistically indistinguishable from zero for volumes, the event-study estimates show no break at the announcement, and restricting the sample to granulated sugar leaves a price effect indistinguishable from zero. Prices in the Southeast rose slightly less than prices elsewhere, not more.

The structural estimates point the same way. Accounting for consumer heterogeneity raises the absolute price coefficient roughly threefold relative to the instrumented logit and yields a median own-price elasticity of -1.44, with cross-price elasticities that spread diversion across many competing brands rather than concentrating it on any one rival. Recovered marginal costs are substantially lower for store brands than for national brands, a median of 1.80 against 3.69 cents per ounce, and they move with retail prices through them 2021–22 period of rising input costs. Median markups and Lerner indices are slightly lower after the announcement than before. The acquired brands, Imperial and Dixie Crystals, show declining absolute markups, and the declines are no larger than those of comparable national

brands over the same period.

The counterfactual decomposition isolates the ownership channel by holding demand fixed. The change from separate to joint pricing raises equilibrium prices by 0.03 percent, with effects on markups, profits, consumer surplus, and welfare of the same order. The effect exceeds one-tenth of one percent in only 23 of 343 markets and never exceeds 0.6 percent in any market. The difference between the pre- and post-announcement cost environments generates larger effects, a 0.44 percent decline in prices and a 2.13 percent decline in markups, so the total effect is almost entirely a cost effect.

Three characteristics of this market explain why a transaction of this prominence leaves so little trace in retail markets. The merging parties' combined retail presence is small, with Imperial and Dixie Crystals accounting for 3.3 percent of national volume and U.S. Sugar's retail brands for less than 0.1 percent. Private label, which accounts for 57 percent of category volume, is priced by retailers whose incentives the transaction does not change. And substitution is spread across many brands and three sugar types, so the demand newly internalized under joint pricing is a small fraction of the demand facing either party.

Several limitations qualify these conclusions. The demand estimates depend on the validity of the price instruments and on the random-coefficients specification. Recovered markups and costs depend on the assumption of static multiproduct Bertrand pricing among the firms that own retail brands; contracting between refiners and retailers is not modeled, so the recovered marginal costs measure the full cost of delivering a product to the shelf rather than a refiner's production cost. Prices are quarterly unit values averaged across stores, which removes the promotional variation to which consumers respond most strongly and therefore attenuates measured price sensitivity; estimated markups should be read as upper bounds and recovered costs as lower bounds. Finally, the transaction closed in November 2022, at the end of the sample period, so the reduced-form estimates measure the response to the announcement and the ensuing litigation rather than to a completed change in control, and the structural exercise is correspondingly prospective.

Subject to these qualifications, the evidence does not support the view that this transaction raised retail prices through the consolidation of brand ownership. The finding carries a broader implication for merger analysis. A transaction that substantially raises concentration among producers may change very little about the market in which consumers actually choose, because what determines the unilateral price effect is not the change in producer concentration but how much of the consumer-facing market the merging firms control and how readily buyers substitute between the products they own. In this case both answers point the same way, and the model attributes essentially none of the observed change in retail outcomes to the change in ownership.

References

Enghin Atalay, Ali Hortacsu, and Chad Syverson. Vertical integration and input flows. *American Economic Review*, 104(4):1120–1148, 2014.

Steven Berry. Estimating discrete-choice models of product differentiation. *RAND Journal of Economics*, 25(2):242–262, 1994.

Steven Berry and Philip Haile. Foundations of demand estimation. In *Handbook of Industrial Organization*. 2021.

Steven Berry, James Levinsohn, and Ariel Pakes. Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890, 1995.

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

Gregory Crawford, Robin Lee, Michael Whinston, and Ali Yurukoglu. The welfare effects of vertical integration in multichannel television markets. *Econometrica*, 86(3):891–954, 2018.

David Genesove and Wallace Mullin. Testing static oligopoly models: Conduct and cost in the sugar industry, 1890–1914. *RAND Journal of Economics*, 29(2):355–377, 1998.

David Genesove and Wallace Mullin. Rules, communication, and collusion: Narrative evidence from the sugar institute case. *American Economic Review*, 91(3):379–398, 2001.

David Genesove and Wallace Mullin. Predation and its rate of return: The sugar industry, 1887–1914. *Review of Economic Studies*, 73(3):699–735, 2006.

Justine Hastings and Richard Gilbert. Market power, vertical integration and the wholesale price of gasoline. *Journal of Industrial Economics*, 53(4):469–492, 2005.

Nathan Miller and Matthew Weinberg. Understanding the price effects of the millercoors joint venture. *Econometrica*, 85(6):1763–1791, 2017.

Aviv Nevo. A practitioner's guide to estimation of random-coefficients logit models of demand. *Journal of Economics and Management Strategy*, 9(4):513–548, 2000.

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

Volker Nocke and Michael Whinston. Concentration thresholds for horizontal mergers. *American Economic Review*, 112(6):1915–1948, 2022.

Janusz Ordover, Garth Saloner, and Steven Salop. Equilibrium vertical foreclosure. *American Economic Review*, 80(1):127–142, 1990.

Craig Peters. Evaluating the performance of merger simulation: Evidence from the u.s. airline industry. *Journal of Law and Economics*, 49(2):627–649, 2006.

Patrick Rey and Jean Tirole. A primer on foreclosure. In *Handbook of Industrial Organization*. 2007.

Joseph Spengler. Vertical integration and antitrust policy. *Journal of Political Economy*, 58 (4):347–352, 1950.

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

Matthew Weinberg and Daniel Hosken. Evidence on the accuracy of merger simulations. *Review of Economics and Statistics*, 95(5):1584–1600, 2013.