

Manufacturer-Retailer Relationships and the Distribution of New Products

Bowen Luo*

February 20, 2023

Abstract

Industry practitioners often emphasize the importance of manufacturer-retailer relationships in product distribution. In this paper, I formalize the industry’s notion of a relationship with a repeated game framework and study its impact in the heavily regulated hard cider market. I present evidence that retailers and leading manufacturers coordinate with and offer preferential treatments to each other when setting assortments and wholesale prices. Based on this evidence, I develop a repeated game-based model to estimate each pair’s coordination, which is linked to the manufacturer’s performance at the retailer in the broader beer market. The results show the relationships increase new cider availability by 17.5% and 5.1%, respectively, for Anheuser-Busch InBev and MillerCoors, the two leading brewers. The relationship’s effect is determined jointly by the degree of assortment distortion and the reduction in double marginalization. Although these relationships could improve welfare, they imply that current regulations in the alcoholic beverage industry fail to generate a level playing field for every manufacturer.

Keywords: product distribution, product assortment, vertical relationships, manufacturer-retailer relationships, repeated games, alcoholic beverage markets

*University of Houston. bluo4@central.uh.edu. I am deeply grateful to my PhD advisors Paul Ellickson and Mitch Lovett for their guidance and support. I also thank Tomomichi Amano, Kristina Brecko, Jacob Burgdorf, Hana Choi, Ying Fan, Ron Goettler, Avery Haviv, Sylvia Hristakeva, Yufeng Huang, Jeanine Miklos-Thal, Stephan Seiler, Greg Shaffer, Adam Smith, Takeaki Sunada, Sriram Venkataraman, Chenyu Yang, the four reviewers for the AMS Mary Kay and ISBM dissertation proposal award, and the conference and seminar participants at Bocconi, CKGSB, Columbia, CUHK, Georgia State, HKU, Imperial College London, ITAM, NUS, Rochester, Rotterdam, Tilburg, UCL, UCLA, AMS Conference 2021, ASSA Annual Meeting 2022, 2023, IIOC 2022, ISBM Webinar, and Marketing Science Conference 2021 for their comments. I own analyses calculated (or derived) based in part on (i) retail measurement/consumer data from Nielsen Consumer LLC (“NielsenIQ”); (ii) media data from The Nielsen Company (US), LLC (“Nielsen Media”); and (iii) marketing databases provided through the respective NielsenIQ and the Nielsen Media Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are my own and do not reflect the views of NielsenIQ or Nielsen Media. Neither NielsenIQ nor Nielsen Media is responsible for, had any role in, or was involved in analyzing and preparing the results reported herein. All errors are my own.

1 Introduction

Manufacturers and retailers often rate their relationships as having high importance for their business, especially in determining product distribution. For example, speaking on the condition of anonymity, a leading national manufacturer once said in an interview, “*we have such good relationships with our retailers that we’re able to have really great conversations and get more placement on shelves.*” Anecdotally, retailers hold regular meetings with leading manufacturers and allocate shelf space based on their assortment recommendations. Their distribution advantage can even be carried over to new and adjacent categories, potentially contributing to, for example, Frito Lay’s strong expansion to tortilla chips with Doritos and Pepsi’s expansion to tea drinks with Pure Leaf.

In this paper, I study the impact of manufacturer-retailer relationships on new product markets. Despite their managerial relevance, such relationships have not been well formalized in the vertical relationship literature, and empirical evidence from field data is scarce. Using a repeated game framework, I formalize the industry’s notion of a relationship—if a manufacturer-retailer pair has a good relationship and high trust in each other, they will have a high coordination level in their repeated game. Specifically, I focus on two cooperation channels in the relationship: when a retailer gives preferential treatment to a manufacturer when setting assortments, and when in return, the manufacturer gives preferential wholesale discounts to the retailer. I present empirical evidence of both channels and develop a model to estimate their impact.

While a good relationship can lead to a double win for both firms, it can also reduce the availability of other manufacturers’ products, even in a heavily regulated environment. In fact, foreclosing rivals’ products is one of the most important vertical issues industry regulators evaluate. Most policy discussions and empirical papers focus on contractual arrangements (e.g., Conlon and Mortimer, 2021) and fixed transfers (e.g., Hristakeva, 2022b,a; Sudhir and Rao, 2006). However, the relationships studied here can also bias assortments towards leading manufacturers. Moreover, they are hard to regulate—even if the regulations restrict explicit contracts and ban fixed transfers, a pair can use an implicit relational contract to distort assortments. Therefore, manufacturer-retailer relationships can reduce the ability of vertical regulations to generate a level playing field for every manufacturer.

Because separating these relationships from other vertical arrangements is empirically challenging with typical retail scanner data, I exploit a unique setting in the US hard cider market to study them.

This setting has two advantages. First, hard cider was a new market with new brands introduced by big brewers, such as Anheuser-Busch InBev (ABI) and MillerCoors. Because the cider and beer brands are not related, and because lager beer and hard cider cater to different consumer segments, I use a brewer's share in the beer category at a given retailer before the cider launched to proxy their long-term relationship. This cross-sectional, long-term measure reflects the trust the two firms have built up through past collaborations, as well as their anticipation of future interactions. I then explore how the cider assortments and prices are linked to this measure of manufacturer-retailer relationships. Second, the heavy regulations in the alcoholic beverage industry can help identify the relationship's effect and rule out alternative explanations. Specifically, federal law (Title 27 of the USC and CFR) prohibits offering things of value to retailers (e.g., slotting fees and services) and forcing tie-in sales (i.e., requiring a retailer to buy one product in order to buy another). Also, the wholesale pricing regulations differ across states, which provides variations to pin down preferential wholesale discounts in states where cider wholesale prices are not regulated.

I begin with rich descriptive evidence that supports the strong impact of manufacturer-retailer relationships on new product distribution. These relationships are established between retailers and the leading brewers, ABI and MillerCoors. Across manufacturers, I find the new ciders produced by ABI and MillerCoors are available in more stores than other ciders with similar sales performance. Across retailers, retailers that sell a lot of ABI's beer tend to carry its new cider, and the same goes for MillerCoors. Across manufacturer-retailer pairs, I find a strong and positive association between a brewer's past beer share at a retailer and the retailer's adoption of the brewer's new cider. I interpret this result as a relationship effect, and I use industry regulations and high-dimensional fixed effects to rule out alternative explanations, including slotting fees, local product preferences, and synergy in transportation and distribution. The relationship's effect on distribution is further confirmed by additional evidence that exploits ownership changes of incumbent cider brands.

I also find that the brewers offer wholesale discounts on cider to their preferred retailers in states where cider wholesale prices are not regulated. Because these retailers are mostly multi-state retailers that operate in both regulated and unregulated states, the wholesale discounts serve to compensate retailers for their preferential cider distribution in all states. Although wholesale prices are not directly observed, I infer the preferential wholesale discounts by comparing the correlation between a brewer's past beer share and its cider retail price at a given retailer across the regulated

and unregulated states. I find a negative correlation between the two in the unregulated states but not in the regulated states, suggesting the manufacturers offer wholesale discounts to their preferred retailers in unregulated states. This negative correlation is the opposite of what a wholesale price bargaining model would predict (e.g., Ellickson et al., 2018; Draganska et al., 2010; Gross, 2019), but is consistent with a repeated game model that also takes into account preferential assortments.

I then develop a repeated game framework to model the relationships and coordination between manufacturers and retailers. Each period of a repeated game has an assortment stage, a wholesale pricing stage, a standard retail pricing stage, and a demand stage. A retailer can give preferential treatment to a manufacturer when setting assortments, while the manufacturer can offer preferential discounts to the retailer when setting wholesale prices. Their coordination is subject to the incentive compatibility constraints, where deviation is punished by reverting to the static Nash equilibrium. Importantly, relationship trust is captured by the probability a pair believes their relationship will continue in the future, and it determines the coordination level they can achieve.

Because it would be formidable to simultaneously estimate repeated games for many manufacturer-retailer pairs, I derive a tractable empirical model from the underlying repeated games based on Fan and Sullivan (2018). Fan and Sullivan (2018) provide a micro-founded empirical model consistent with supergames in a horizontal setting. I adapt their model and add additional assumptions to address the differences between horizontal and vertical settings. The derived static model for a repeated game has three terms in each player's objective: the player's own profit, the partner's profit, and the partner's deviating profit. The parameter of the final term is interpretable as the (inverse of the) coordination level, and it depends on the manufacturer's past beer share at the retailer. The parameters are identified by demand shifters and differences in state regulation for the wholesale pricing stage, and by how assortments vary with profits across markets for the assortment stage.

The structural model results show that manufacturer-retailer relationships play a significant role in both retail assortment and wholesale pricing decisions. Relationship coordination increases with the past beer share. Overall, the estimates reveal moderate but imperfect coordination between retailers and the two leading brewers, ABI and MillerCoors, which together account for more than half of the beer market and have been doing business with retailers for many decades. They offer average wholesale discounts of 7.6% and 5.5%, respectively, in states where wholesale prices are not explicitly regulated.

Using the model estimates, I quantify the impact of manufacturer-retailer relationships on distribution, profits, and welfare. The relationships increase cider availability by 17.5% and 5.1% and cider profits by 9.7% (\$0.5m) and 1% (\$0.1m) for ABI and MillerCoors, respectively. Retailers' profits increase by 1.1% (\$2.1m), suggesting their gain from wholesale discounts outweighs their loss from sub-optimal assortments. The leading manufacturers and retailers' bigger pie comes from the increased availability of the manufacturers' products and the reduction in double marginalization. This double win is at the expense of 1.9% (-\$2.1m) profits for the remaining manufacturers due to their lower availability. Consumers see a welfare increase of 1.1% (\$2.3m) because of the reduced prices, despite the worse assortments. Although these relationships are hard to regulate, implementing posted wholesale prices in the currently unregulated states could break them down by preventing brewers from compensating their preferred retailers. Thus, the counterfactual results could also inform the impact of a nationwide tightened wholesale pricing law, which would lead to higher prices in the currently unregulated states and better assortments across all states.

The paper has several important policy implications for vertical regulations. First, it demonstrates that manufacturer-retailer relationships bias assortments towards the leading manufacturers and reduce the availability of small manufacturers' products. Moreover, because these relationships rely on implicit agreements and are hard to regulate, they restrict regulators' ability to ensure product variety and restore a fair competitive environment. While the regulatory discussion around vertical relationships largely focuses on exclusive contracts and slotting fees, this paper shows that banning both elements would prompt manufacturer-retailer pairs to leverage relationships and would still bias assortments. Nevertheless, regulators should still ban slotting fees, because that would force manufacturers to compensate retailers with wholesale discounts, which can reduce double marginalization and improve welfare. To actually break down these relationships, regulators need to cut down the channels the pairs secretly exchange benefits. In this case, implementing posted wholesale prices across all states would break them down by preventing brewers from compensating their favored retailers.

Furthermore, the alcohol industry is the focus of some recent policy discussions on promoting competition in the US economy.¹ Regulators are concerned about the exclusionary trade practices

¹See Sections 5(j) and 5(k) of a July 2021 Executive Order <https://www.whitehouse.gov/briefing-room/presidential-actions/2021/07/09/executive-order-on-promoting-competition-in-the-american-economy/> and a follow-up report by the Department of the Treasury's Alcohol and Tobacco Tax and Trade Bureau (TTB)

and entry barriers faced by small producers in this industry. This paper provides evidence that the current regulatory efforts in the alcoholic beverage industry do not successfully generate a level playing field for every manufacturer. Regulators should thus factor in pair relationships when updating the regulations for alcoholic beverages.

Broadly, these relationships can contribute to leading manufacturers' persistent dominance and the high concentration in consumer packaged goods (CPG) markets. Even if their main products experience stagnation, leading manufacturers have a competitive advantage in new and adjacent product categories thanks to their relationships with retailers. Thus, the relationships serve as an important channel through which leading manufacturers maintain their dominance and thereby contribute to long-term CPG market concentration.

This paper contributes to the literature on retail vertical relationships. The closest paper is Hristakeva (2022b), who proposes an estimation strategy to examine equilibrium market outcomes under different vertical contract forms. She shows that lump-sum transfers affect not only product assortments but also potentially prices through replacement threats. In contrast, my paper presents empirical evidence for the impact of manufacturer-retailer relationships through preferential product assortments and wholesale discounts, where the latter directly reduce retail prices and double marginalization. Moreover, my paper suggests that banning fixed transfers would prompt manufacturer-retailer pairs to leverage relationships instead of creating a competitive environment. This paper also differs from papers on category captaincy (e.g., Viswanathan et al., 2020) in that the relationships are non-contractual, do not necessarily involve services, and can exist between retailers and multiple manufacturers. Finally, this paper adds to the broader literature on retail product adoption (e.g., Bronnenberg and Mela, 2004; Hwang et al., 2010; Misra, 2008) and relationship marketing (e.g., Dwyer et al., 1987; Heide, 1994; Weitz and Jap, 1995; Kaufman et al., 2006; Zhang et al., 2014, 2016), which typically uses survey data to measure relationships.

The rest of the paper is organized as follows. Section 2 describes the setting, regulations, and data. Section 3 presents the descriptive evidence of manufacturer-retailer relationships' effects on cider distribution and wholesale discounts. Section 4 provides a structural model of these relationships. Model estimation and identification are discussed in Section 5, then Section 6 presents the results. Section 7 presents the counterfactual analysis that quantifies the relationship's impact.

<https://home.treasury.gov/system/files/136/Competition-Report.pdf>

Section 8 concludes the paper.

2 Industry Background and Data

This section first provides the industry background of hard cider and discusses alcoholic beverage regulations. It then describes the data for analysis and presents important patterns in the data.

2.1 US Hard Cider Market

Hard cider saw a nine-fold growth in the early 2010s—the fastest-growing segment in the US alcoholic beverage industry—fueled by consumer interest in gluten-free, healthy products.² By 2018, its sales reached \$1.2 billion, including both on- and off-premise sales. Hard cider is made from fermented apple juice, with ABV (alcohol by volume) typically ranging from 4.5% to 7%. While some niche, high-ABV cider (up to 12% ABV) caters to wine drinkers, most cider products are carbonated and designed as refreshing, beer-like drinks. The latter is the focus of this paper. Like craft beer, hard cider is considered a premium alcoholic drink, but it has a sweeter taste, is gluten-free, and caters to a younger and gender-balanced profile. It is usually sold next to craft beer in stores.

To participate in the growth of hard cider, major brewing companies introduced new brands, including Angry Orchard (by Boston Beer in 2012), Smith & Forge (by MillerCoors in 2014), and Johnny Appleseed (by ABI in 2014). In addition, they acquired incumbent cider brands (e.g., Crispin’s acquisition by MillerCoors). In the meantime, retailers expanded their cider assortments to meet increasing demand. As shelf space is limited, retailers have to choose which new cider to adopt, and manufacturer-retailer relationships can play an important role in these decisions.

2.2 Industry Regulations

Cider, like other alcoholic beverages, is heavily regulated in the US. Policies are established at the federal and state level to tax and regulate the sale and distribution of alcoholic beverages.³ The goal of these policies is to encourage fair competition and maintain retailer independence. As a result,

²Source: IBISWorld Industry Report and Nielsen Report https://ciderassociation.org/wp-content/uploads/2019/02/Nielsen-Presn-at-CiderCon-2019_2-7-2019.pdf.

³Some local governments establish additional regulations for alcoholic beverages in their jurisdictions, which are not considered in this paper.

store brands are rare, and manufacturers have few tools to promote their products. In particular, they need to follow restrictions on trade and wholesale pricing practices.

In general, federal law prohibits industry members (including manufacturers and wholesalers) from offering things of value to retailers for better product placement, distribution, and/or advertising (Gundlach and Bloom, 1998).⁴ Here, things of value include slotting fees, trade money, and services. Category management is included in services of value. Tie-in sales, defined as an industry member requiring a retailer to purchase one product in order to obtain another product, are also not permitted. These policies are actively enforced by the Alcohol and Tobacco Tax and Trade Bureau (TTB), because the practices may result in the exclusion of competitors' products. For example, TTB shut down Kroger's plan to implement slotting fees for its beer category in 2016.⁵

Wholesale pricing regulations are established by states, which have broad power to regulate alcohol sale and distribution within their borders. The regulations fall into two types: Post & Hold and volume discount restrictions. Post & Hold requires wholesale prices to be posted and held for a certain period so that all retailers have the chance to buy at the same prices. Volume discount restrictions require the same wholesale price per unit to be charged regardless of the quantity purchased. Based on the alcoholic beverage classification in each state, 16 states require Post & Hold (or just Post) for cider and 14 states ban volume discounts, with 10 states overlapping. Since wholesale discounts are feasible only in states without these restrictions, the difference in price patterns between restricted and non-restricted states provides information on wholesale discounts in non-restricted states. More details of the state regulations can be found in Web Appendix A.

Although most states have a "three-tier system" that requires the separate operation of manufacturers, wholesalers, and retailers, wholesalers can be seen as representing the interests of manufacturers and are not explicitly modeled in this paper. Wholesalers do not have much power compared to the other two tiers (Asker, 2016). They typically serve one main account (ABI or MillerCoors) along with some craft brands. In addition, the wholesale market is fragmented and has more than 2,000 wholesalers in total, according to the National Beer Wholesalers Association, and ABI alone

⁴The relevant policies are established by the Federal Alcohol Administration (FAA) Act (Title 27 of the US Code) and Title 27 of the Code of Federal Regulations (CFR). The FAA Act prohibits the following four categories of trade practices: exclusive outlet, tied house, commercial bribery, and consignment sales. The law applies to producers of distilled spirits and wine. It also applies to producers of malt beverages if there is a similar state law, and 48 states have similar trade practice laws.

⁵See the TTB website for a full list of recent enforcement actions.

has 500. Wholesalers are typically assigned exclusive territories by a manufacturer such that the same wholesaler serves all stores in its geographic region for the manufacturer. Thus, a manufacturer’s distribution quality should be the same within a region, and manufacturer-location fixed effects should absorb its variation. On a separate note, this “three-tier system” also prohibits private labels because a retailer cannot operate as a manufacturer.

In summary, first, federal law prohibits offering things of value to retailers, including slotting fees, services, and payments for advertising and display. Although common in other categories, these practices do not affect the distribution of cider. Second, wholesale pricing regulations differ across states, which can be used to examine wholesale discounts. Third, wholesalers can be seen as representing the interests of manufacturers and are not explicitly modeled in the paper. Fourth, the exclusive territory feature implies that manufacturer-location fixed effects can capture the variation in distribution network quality.

2.3 Data

The main data used comes from the Kilts Center’s Nielsen Retail Scanner Data. It contains weekly UPC-level store sales and prices for hard cider from 2011 to 2016. I infer product distribution from store sales, and a store is seen as carrying a product in a given period if the quantity sold is greater than zero (i.e., conditional on the product having any sales). A limitation of this definition is that if a store carries a product but does not sell any in a period, it would be seen as not carrying the product in that period. To mitigate this concern, in the descriptive analysis, I fill in potentially missing carrying observations if a product is sold both in the three weeks before and the three weeks after the focal week. In the structural estimation, I aggregate the data to the monthly level. Because marketing supports could affect new product demand, I use the feature and display information (available for 16% of the stores)⁶ and the brand-level advertising exposures (from the Nielsen AdIntel Data) in the demand estimation. In addition, I manually collect cider product attributes from their packages, nutrition labels, and websites. Relevant attributes include alcohol content, sugar content (correlated with calories), and whether the product is seasonal.

⁶Although the selection criteria of these stores are unknown, I have checked 1) they cover almost all retail chains, i.e., the selection is on stores, not on chains, 2) the average prices and shares of the products are very similar between those with the information and those without for each chain, and 3) descriptive analysis results are close when the sample is a subset of those with the information.

Table 1: Manufacturer Beer Share by Retailer, 2010

Statistic	N	Mean	St. Dev.	Min	Median	Max
Anheuser-Busch InBev	61	0.36	0.12	0.12	0.35	0.59
Boston Beer	61	0.02	0.02	0.001	0.02	0.08
Heineken	61	0.07	0.04	0.01	0.06	0.18
MillerCoors	61	0.27	0.09	0.11	0.25	0.57

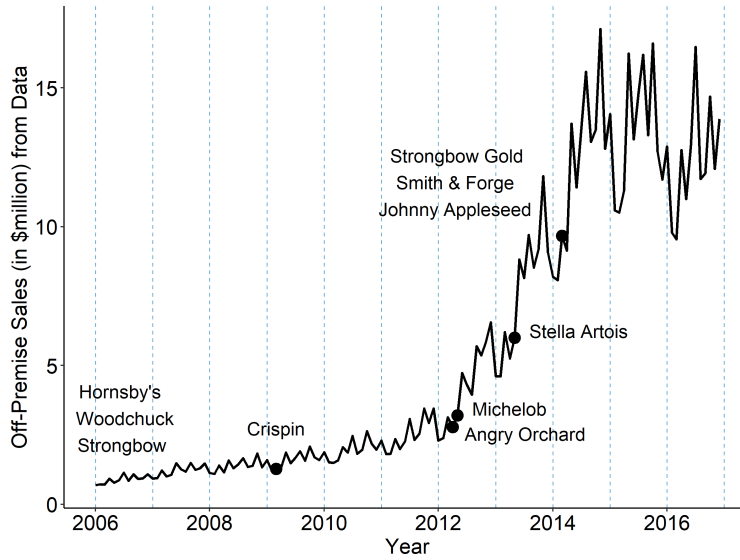
I use beer sales data to construct a proxy for manufacturer-retailer relationship. The measure is the pair-specific manufacturers’ beer shares at retailers in 2010 (the year before the cider data sample). It is effectively a long-run measure because the shares are persistent, with a yearly average serial correlation of 0.996. This measure is a good proxy for a relationship because retailers often communicate with major suppliers and a pair can build trust through past collaborations. When trust is high, each party is more likely to believe the other party would cooperate and the relationship can continue in the future. The measure also reflects possible future interactions. On the other hand, the beer share is an outcome of both the relationship itself and the persistent tastes of the retailer’s customers for lager beer. These tastes motivate the pair to collaborate and maintain a good relationship. Table 1 presents the summary statistics for the measure, which shows large variation across retailers for a given manufacturer’s shares. As the leading brewers, ABI and MillerCoors together account for more than half of the beer market. Note that hard cider remains a small category relative to beer ($< 2\%$) throughout the sample. I use the beer shares in 2010 to avoid the direct effect of contemporaneous beer shares and the reverse effect of the hard cider launches on beer shares. I call this measure “past beer share” henceforth.

I collect federal-level and state-level regulations for alcoholic beverages. The federal regulations on trade practices are drawn from the Federal Alcohol Administration Act (Title 27 of the US Code), Title 27 of the Code of Federal Regulations, and TTB. The state regulations on wholesale pricing practices are collected from the National Institute on Alcohol Abuse and Alcoholism, as well as the state statutes and regulations.

2.4 Summary Statistics

Figure 1 plots monthly cider sales in the data from 2006 to 2016. As we can see, the category gradually expanded before 2010 and took off around 2012. After 2012, six new brands were launched

Figure 1: Hard Cider Sales and Launch Timing



Notes: Dots indicate the time the brands were launched. The three brands on the left were in the market before the sample starts.

nationally by big brewers, which fundamentally changed the cider market structure. The market reached its peak around 2015 and remained stable after that. Cider has a strong seasonality, with more sales in summer and fall and less in winter.

Table 2 provides a summary of dollar shares, distribution, and prices of major cider brands. The table has two panels. The upper panel lists the new cider brands launched by brewers, while the lower panel lists the top incumbent brands before the launches. Columns (3) and (4) present two measures of cider distribution—the percentage of stores that carry the cider brand (% Stores) and the percentage of stores that carry the cider brand weighted by store beer sales (% PCV, product category volume). Column (5) shows the price of a 6-pack (72oz) of cider. We can see the price positioning is similar across cider brands and close to craft beer. This table is based on statistics from the last quarter of 2014 when all the expansion, launches, and acquisitions were complete and the market’s stable period began.⁷

Table 2 provides several key observations of the market.⁸ First, large brewers show great interest in this growing category, and the market is concentrated. The top five manufacturers (and importers)

⁷Some changes since 2015: Johnny Appleseed and Hornsby’s exited the market in the second half of 2016. Regional craft brands (e.g., Bold Rock, 2 Towns Ciderhouse, Austin Eastciders) have been growing in recent years.

⁸Angry Orchard is the only brand that used test markets (selling for a few months in New England) before their national launch. The other brands were launched nationally without test marketing. I ignore the test market period of Angry Orchard in my analysis.

Table 2: Summary Statistics


Brand (1)	Share (2)	% Stores (3)	% PCV (4)	Price (5)	Manufacturer/Importer (6)
<i>New Brands</i>					
Angry orchard	54.8%	70.4%	95.1%	\$8.82	Boston Beer
Strongbow gold	5.3%	30.6%	61.8%	\$9.28	Heineken USA
Smith forge	5.1%	44.3%	76.3%	\$8.62	MillerCoors
Johnny appleseed	4%	46.9%	81.2%	\$8.37	Anheuser-Busch InBev
Stella artois	3.5%	32%	62.7%	\$10.82	Anheuser-Busch InBev
Michelob	1.5%	19.9%	39.3%	\$7.68	Anheuser-Busch InBev
<i>Incumbent Brands</i>					
Woodchuck	10%	37.4%	69.4%	\$8.74	C&C Group PLC
Crispin	3.7%	23.7%	51.5%	\$10.53	MillerCoors
Hornsby's	1.7%	14.2%	33.4%	\$8.08	C&C Group PLC
Strongbow	1.6%	12.5%	25.4%	\$9.33	Heineken USA

account for 91.3% of the market share, as we can see from the dollar shares in column (2) and the manufacturers in column (6). Angry Orchard from Boston Beer is the absolute leader and takes half of the market. Most brands listed in the table have fairly high market penetration (columns (3) and (4)), even though their shares are much smaller than Angry Orchard's shares. Large stores (in terms of beer sales) carry more cider brands than small stores, as evidenced by the higher % PCV than % Stores. As the cider category grows, the top incumbent brands become the acquisition targets of big companies—all four incumbent brands have changed hands to big companies.

More importantly, most of the cider brands are new and not associated with their parent brewers' signature lagers, except for Stella Artois and Michelob. This is because cider is on the premium side, and the brewers might think their lager brands have a non-crafty image. Moreover, the cider brands have separate labels and websites that do not mention their parent brewers, so consumers are unlikely to recognize the association. Thus, umbrella brand spillover is minimal in this market and has little impact on cider distribution.

Lastly, the cider from ABI and MillerCoors is available in more stores than other cider with similar or better sales performance. For example, more stores carry Johnny Appleseed (46.9%) and Smith & Forge (44.3%) than Strongbow Gold (30.6%) and Woodchuck (37.4%), and the same goes for Michelob (19.9%) versus Hornsby's (14.2%) and Strongbow (12.5%). In other words, despite lower shares, ABI and MillerCoors have better cider distribution than other manufacturers.

Figure 2: R^2 comparison for product distribution, quantities, and prices



Note: The figure is based on the six new brands introduced by major brewers from 2012 to 2014.

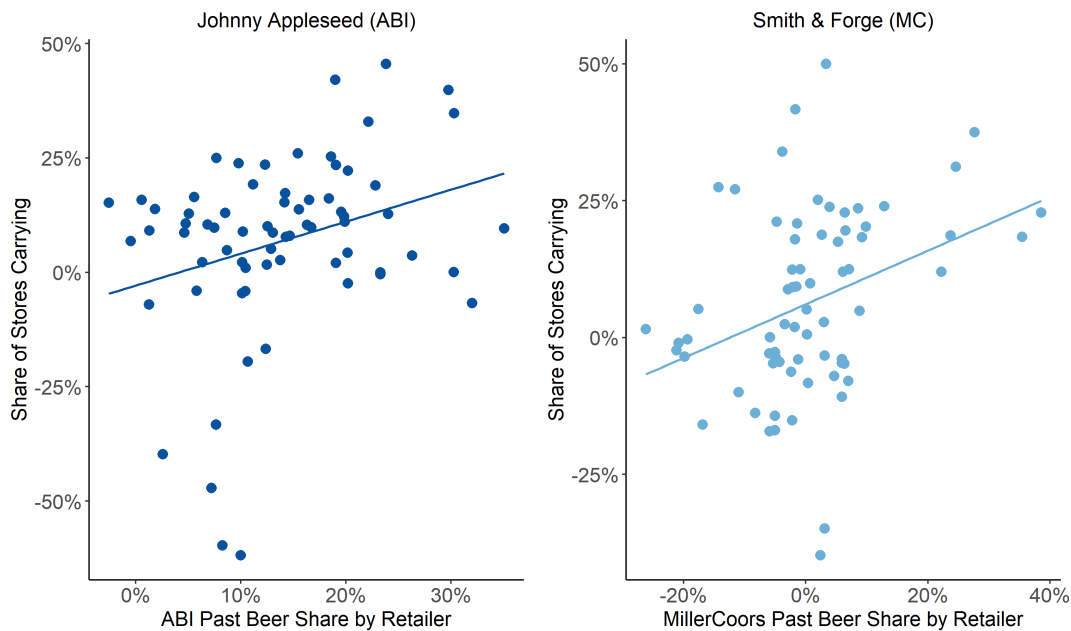
2.5 The Importance of Pair Relationships in Cider Distribution

Beyond the across-brand differences described above, pair relationships have a large explanatory power in cider distribution. Figure 2 presents the R-squared comparison of how much different fixed effects explain the remaining variations in product distribution, quantities, and prices, beyond the product-week fixed effects. It shows product-chain(retailer)-week fixed effects explain 43% of the remaining variation in product distribution, while product-county-week fixed effects can only explain 20%. Moreover, the opposite is true for quantities and prices: product-county-week fixed effects explain more variation than product-chain-week fixed effects. Therefore, pair relationships have an important role in cider distribution, one that is not explained away by demand factors. The lower uniformity of prices relative to distribution within a chain can be explained by the difference in excise tax and wholesale pricing regulations across states.

3 Descriptive Evidence

In this section, I present descriptive evidence of manufacturer-retailer relationships' effects on new cider distribution and wholesale discounts. I start with data plots and fixed effects regressions to show a strong and positive association between a manufacturer's past beer share at a given retailer and that retailer's adoption of the manufacturer's new cider. I interpret the result as a relationship effect, and I rule out slotting fees and other explanations using a combination of industry regulations

Figure 3: New Cider Distribution and Manufacturer Past Beer Share across Retailers



Notes: The plot shows the percentage of stores that carried the cider across retailers in October, 2014 against the manufacturer’s beer share by retailer in 2010. The negative numbers are due to retailer fixed effects (across brands) being removed.

and high-dimensional fixed effects. I then exploit the difference in wholesale pricing regulations across states to show that brewers offer wholesale discounts to preferred retailers in states where cider wholesale prices are not explicitly regulated. Finally, as additional evidence, I show the same distribution pattern applies to craft cider brands that were acquired by major brewers.

3.1 New Cider Distribution and Manufacturer Past Beer Shares

Plots. If relationships with retailers can increase a manufacturer’s new product distribution, the manufacturer’s past beer shares at retailers should positively correlate with the pair-specific new cider adoption. Figure 3 plots the availability of Johnny Appleseed (introduced by ABI) and Smith & Forge (introduced by MillerCoors) across retailers—the x -axis is the past beer share, and the y -axis is the share of stores that carry the cider brands. Because hard ciders are generally more popular at retailers that sell less domestic lager beer, I partial out retailer fixed effects to remove the category-level popularity. The plots use data from October 2014 and reflect the cider distribution in stable periods.

As Figure 3 shows, the percentage of stores carrying a new cider positively correlates with the

manufacturer’s past beer share at the retailer. In other words, heavy Budweiser sellers tend to carry ABI’s Johnny Appleseed, while heavy Miller and Coors sellers tend to carry MillerCoors’ Smith & Forge. A 10 percentage-point increase in the brewers’ past beer shares corresponds to a 7.0 (left) and 4.9 (right) percentage-point increase in the share of stores that carry the cider, and both slopes are statistically different from zero. This pattern is consistent with the story that good pair relationships increase new product distribution.

Regression analysis. To show the pattern more formally and to control for more confounds, I conduct fixed effects regressions of new cider distribution on past beer shares:

$$cider_distribution_{bst} = beer_share_{m(b)r(s)}^{2010} \times \beta_1 + s_{br(s)}^{u,2010} \times \beta_2 + \lambda_{bc(s)t} + \lambda_{st} + \epsilon_{bst}, \quad (1)$$

where $beer_share_{m(b)r(s)}^{2010}$ is the past beer share measure—the beer share of manufacturer m at retailer r in 2010, $s_{br(s)}^{u,2010}$ is brand b ’s beer share at r in 2010 if b is an umbrella brand (i.e., Stella Artois or Michelob), and $\lambda_{bc(s)t}$ and λ_{st} are brand-county-week and store-week interactive fixed effects. I use two variables to measure product distribution: a dummy that equals one if store s carries brand b in week t , and the log number of b ’s products carried by s in week t , conditional on s carrying b . The first variable measures brand availability and the “extensive margin,” while the second variable measures brand presence and the “intensive margin.” The first one is more important because it determines whether consumers of a store have access to a brand at all. I focus on the new brands introduced by brewers and the period they were launched (2012-2014).

The coefficient of interest β_1 captures the extent to which brewers’ past beer shares across retailers predict their new cider distribution. Because the past beer share variable is time-invariant, the variation to estimate β_1 is cross-sectional, across-pair variation instead of within-pair, across-time variation. Because λ_{bct} absorbs the across-brand variation and λ_{st} absorbs the across-retailer variation, the variation to estimate β_1 is within-brand, across-retailer and within-retailer, across-brand. In addition, λ_{bct} controls for 1) transportation costs and wholesaler-specific factors (due to exclusive territories), 2) local consumer tastes and TV advertising (which is at the brand-DMA-week level), and 3) seasonality and brand-specific growth. λ_{st} controls for store-specific changes in cider popularity and expansion of cider shelves. Note that the brand-retailer fixed effects are not included in the regression, because they would otherwise absorb all variation in the past beer share variable.

Table 3: New Cider Distribution and Manufacturer Past Beer Share by Retailer

	carry	carry	carry	log(#UPC)
	(1)	(2)	(3)	(4)
Mfr. Past Beer Share by Retailer	0.44*** (0.17)	0.50*** (0.15)	0.47*** (0.15)	1.53*** (0.53)
Umbrella Brand Share by Retailer		6.75*** (1.22)	5.65*** (1.11)	-3.24 (2.29)
Brand-County-Week FE	Yes		Yes	Yes
Brand-Week FE		Yes		
Store-Week FE	Yes	Yes	Yes	Yes
Observations	10,499,897	10,499,897	10,499,897	3,320,277
R ²	0.70	0.64	0.70	0.88

Notes: Standard errors are clustered at the brand-chain level. *p<0.1; **p<0.05; ***p<0.01

I include the two umbrella brands and use $s_{br(s)}^{u,2010}$ to control for their umbrella brand popularity instead of dropping them, because the main variable $beer_share_{m(b)r(s)}^{2010}$ is mostly driven by the brewer’s other, more popular beer brands (e.g., Budweiser).

Results and discussion. Table 3 presents the regression results of the positive association between past beer shares and new cider distribution. Columns (1)-(3) report the results for the binary brand availability measure “carry,” with column (3) as the preferred model. In the preferred model, a 10 percentage-point increase in the past beer share is associated with a 4.7 percentage-point increase in cider brand availability, confirming the pattern shown in Figure 3. The result is robust to not including the umbrella brand share (column (1)) and to using brand-week fixed effects instead of brand-county-week fixed effects (column (2)). The latter implies that the brand-county factors, such as the brewers’ distances to stores, are not the key confounders of the results. The brand presence of brewers’ new ciders is also positively associated with past beer shares, as shown in column (4). A 10 percentage-point increase in the past beer share is associated with a 15.3% increase in the number of cider products being carried by the retailer, conditional on the retailer carrying the cider brand.

This positive association between past beer shares and new cider distribution suggests a strong impact of manufacturer-retailer relationships. As discussed in Section 2.3, the past beer share is a good proxy for a relationship, because a manufacturer-retailer pair can build up trust through past communication and collaborations. The share also reflects their anticipation of future interactions. In

addition, the fixed effects and industry regulations rule out many alternative explanations, including slotting fees, tie-in sales, and transportation costs.⁹ These results demonstrate a strong cross-category spillover of manufacturer-retailer relationships and highlight the competitive advantage of leading manufacturers in the vertical channel.

The Web Appendix presents several robustness checks and rules out more confounding factors. In Web Appendix B.3, I show that a brewer’s current beer share does not drop when a retailer carries its cider, suggesting the cider distribution result is not driven by a direct replacement of the brewer’s beer on the shelf.¹⁰ Web Appendix B.5 presents a test that drops the highest share brewer of each retailer, and I find the distribution result is not driven by category captaincy. In addition, captains in the alcoholic beverages markets are different from typical captains in that they cannot provide free services to retailers, even though they can still maintain a close relationship with them and offer assortment recommendations. Lastly, in Web Appendix B.6, I show that relationships are established between manufacturers and retailers, instead of between manufacturers and stores, by including the brewer’s past beer share at a given store in the regression and comparing the coefficients.

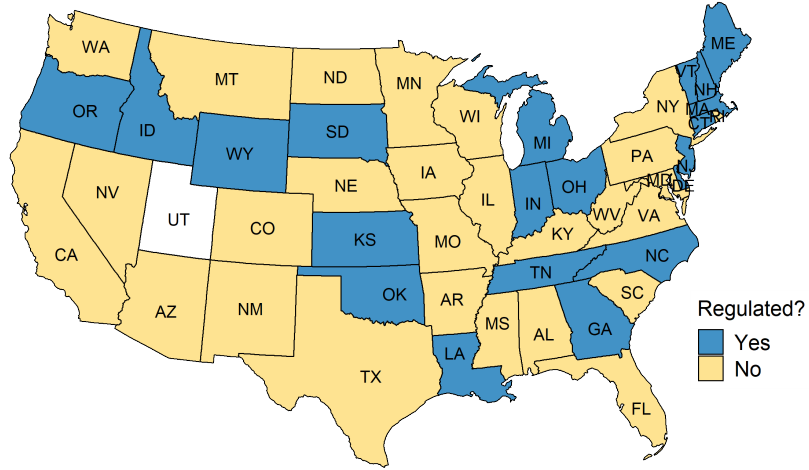
3.2 Wholesale Discounts and Manufacturer Past Beer Shares

Leading brewers have limited ways to compensate retailers for this extra distribution due to the restrictions on fixed transfers and other vertical practices. One possible way is to offer wholesale discounts to preferred retailers in states where they are allowed to do so. Because states have different regulations on cider wholesale pricing, this difference can be exploited to examine the association between manufacturer-retailer relationships and wholesale conduct. I call the states that require Post & Hold and/or ban volume discounts “regulated” states, and call the rest “unregulated” states. Because both rules significantly restrict the preferential discounts manufacturers can offer, I place states with either or both rules in the “regulated” category. Figure 4 shows a map of the two types

⁹An efficiency gain in shipping costs is unlikely to explain the results in Table 3 for two additional reasons. First, the focal manufacturers are all large brewers, so their beer need to be shipped to stores regularly anyway, and cider shipping is just a small part of it. Second, the results are robust to only focusing on brand-store pairs with positive manufacturer past beer shares at the store (see Web Appendix B.2), so they are not driven by last-mile shipping efficiency.

¹⁰There is usually a separate shelf for cider next to craft beer in stores, so the focal brewers’ beer is mainly on a different shelf, and direct exchange is very unlikely. Also, most cider products have single facing. One might worry that even if tie-in sales are illegal, brewers could still offer wholesale discounts for their beer to induce retailers to buy their cider. Web Appendix B.4 shows that the distribution results are robust for states where beer wholesale prices are regulated.

Figure 4: Map of Regulated and Unregulated States for Cider Wholesale Pricing



of states.

If manufacturers offer preferential wholesale discounts to retailers in states where they are allowed to do so, there should be a difference in the association of cider retail prices and past beer shares across the two types of states. If cider retail prices are negatively associated with past beer shares in unregulated states, but not associated in regulated states, this suggests preferential *wholesale* discounts are offered in unregulated states. If instead cider retail prices are negatively associated with past beer shares in both types of states, this suggests preferential *retail* discounts are offered in both types of states, instead of the wholesale discounts offered in unregulated states.

To this end, I run a fixed effects regression similar to equation (1) with cider retail prices as the dependent variable. It includes an interaction term of past beer share and a dummy for regulated states to capture the difference across the two types of states. The regression is at the product-store-week level and includes only products that are on the shelf.

$$\log(\text{cider_price}_{jst}) = \text{beer_share}_{m(b)r(s)}^{2010} \times \beta_1 + \text{beer_share}_{m(b)r(s)}^{2010} \times D_s \times \beta'_1 + s_{br(s)}^{u,2010} \times \beta_2 + s_{br(s)}^{u,2010} \times D_s \times \beta'_2 + \lambda_{jc(s)t} + \lambda_{st} + \epsilon_{jst}, \quad (2)$$

where j indexes products, D_s is a dummy for whether store s is in a regulated state, and the other variables follow the same definitions as in equation (1). The coefficient β_1 captures the association of past beer shares and cider retail prices in unregulated states, and $\beta_1 + \beta'_1$ captures this association

Table 4: Retail Price and Manufacturer Past Beer Share by Retailer

	log(Price)	
	(1)	(2)
Mfr. Past Beer Share by Retailer	-0.102** (0.048)	-0.126** (0.053)
× $\mathbf{1}\{\text{P\&H or Ban Q.D.}\}$		0.124** (0.060)
Umbrella Control?	Yes	Yes
Product-County-Week FE	Yes	Yes
Store-Week FE	Yes	Yes
Observations	8,077,969	8,077,969
R ²	0.987	0.987

Notes: Standard errors are clustered at the brand-chain level. *p<0.1; **p<0.05; ***p<0.01

in regulated states.

Table 4 presents the results of the negative association between cider retail prices and past beer shares in unregulated states. The negative association exists only in unregulated states (-0.126), but not in regulated states (-0.002). The difference is statistically significant and suggests that manufacturers offer preferential wholesale discounts to retailers with good relationships in states where they are allowed to do so. These discounts serve as compensation for retailers.

Note that this result for wholesale prices differs from what a bargaining model would imply. A bargaining model would predict relatively high wholesale prices for manufacturer-retailer pairs with high beer shares because beer shares should be positively related to manufacturer bargaining power. However, once product assortments are taken into account, as they are here, persuading a retailer to carry its products is more important for a manufacturer than giving up some margins. Lower wholesale margins coupled with better distribution could be a win-win solution for both the manufacturer and retailer.

In Web Appendix B.7, I show that the relationship's effect on cider distribution exists in both regulated and unregulated states, unlike the effect on wholesale discounts. Because almost all retailers in the sample are multi-state retailers that operate in both types of states, the results suggest that retailers offer preferential product distribution in all states, while manufacturers compensate them in unregulated states.

Table 5: Incumbent Brand Availability, before/after Ownership Change

	Before	After
	(1)	(2)
New Owner's Past Beer Share by Retailer	0.02 (0.30)	0.88*** (0.32)
Brand-County-Week FE	Yes	Yes
Store-Week FE	Yes	Yes
Observations	3,472,499	5,273,256
R ²	0.82	0.78

Notes: The incumbent brands have a single product most of the time. Standard errors are clustered at the brand-chain level. *p<0.1; **p<0.05; ***p<0.01

3.3 Distribution of Incumbent Brands after Acquisitions

As additional evidence and a validity check for the main finding of the relationship's effect on cider distribution, I exploit incumbent cider brands' acquisitions by major brewers. Specifically, Crispin was acquired by MillerCoors in 2012. Strongbow's importation right was transitioned to Heineken USA as of 2013, to facilitate the Irish cider company C&C's acquisition of Vermont Hard Cider Company. The Appendix shows that Crispin and Strongbow experienced significant distribution increases after their acquisitions by major brewers, whereas the two other incumbent brands acquired by C&C did not.

These ownership changes provide within-brand variations to identify the relationship's effect on cider distribution. If a brewer has an established relationship with a retailer, the brewer should be able to increase the acquired cider's distribution at the retailer after the acquisition. Thus, the new owners' past beer shares should be positively associated with the acquired brands' distribution *after* the ownership changes, but not *before*. The latter also rules out the possibility that the brewers selectively acquired ciders that cater to their beer drinkers. I estimate equation (1) for the incumbent brands separately with data before and data after the acquisitions by brewers.

Table 5 presents the results for incumbent brands, and they support the main distribution result. The incumbent brands' availability is positively associated with the new owners' past beer shares after their acquisitions (0.88). By contrast, the association is small and statistically insignificant before the acquisitions (0.02), which serves as a placebo test.

4 Model

In this section, I formalize manufacturer-retailer relationships with repeated games and derive an empirical model for the games. If a manufacturer-retailer pair has a good relationship and high trust in each other, they will have a high coordination level in their repeated game. I use the probability that the pair believes their relationship can continue in the future to capture their trust in each other. To link the model with the descriptive evidence, I specify the coordination level to depend on the manufacturer’s past beer share at the retailer. Also, the pair’s contact in the beer category can relax their incentive constraints in the cider category, in the spirit of Bernheim and Whinston (1990).

While simultaneously estimating repeated games for many manufacturer-retailer pairs would be formidable, I derive a tractable empirical model from the underlying repeated games. The derivation is based on Fan and Sullivan (2018), who derive an empirical static model theoretically founded on an infinitely repeated game in a horizontal pricing setting. I adapt their model and add assumptions to address the differences between horizontal and vertical settings. Following Fan and Sullivan (2018), the parameters of the empirical model are interpretable as the degree of coordination and asymmetry in the relationships. Compared to the reduced-form analysis, the structural model takes into account two important demand-side factors in assortment choices: retailer-specific brand preferences and substitution between similar products. A major goal of the model is to predict the counterfactual demand and prices, had the stores carried different sets of products. Below, I first show the mapping between an infinitely repeated game and the empirical model, then I present the details of each stage of the model (assortment choice, wholesale pricing, retail pricing, and purchase).

4.1 Relation to a Repeated Game

Consider an infinitely repeated game between a manufacturer and a retailer with grim trigger strategies (i.e., indefinite reversion to Nash equilibrium punishment). Let δ be the discount factor and ϕ be the relationship continuation probability, defined as the probability the players expect their repeated game to continue in the next period. Each period has four stages:

- Stage 1 (assortment choice): Retailer chooses assortments
- Stage 2 (wholesale pricing): Manufacturer sets wholesale prices

- Stage 3 (retail pricing): Retailer sets retail prices
- Stage 4 (purchase): Consumers make purchases

This model timeline of the stage game follows Villas-Boas (2007), Hristakeva (2022b), and Fan and Yang (2020). As suggested by the descriptive evidence, manufacturer-retailer relationships play a role in both assortment choice and wholesale pricing. Thus, in stages 1 and 2, the discounted sum of profits $\sum_{t=0}^{\infty} \delta^t \pi_{rt}$ and $\sum_{t=0}^{\infty} \delta^t \pi_{mt}$ are maximized. In stages 3 and 4, the static objectives are maximized. I assume the retailer makes separate decisions for assortments and prices because they are typically controlled by different units of the retailer. In this paper, I only consider manufacturers' cider profits and assume away the cider's impact on their beer profits. Since cider is very small compared to beer, this is a reasonable assumption, and I also show that a brewer's cider does not particularly substitute away its beer, in Web Appendix B.3.

Consider a Pareto optimal equilibrium of the infinitely repeated game. An equilibrium is Pareto optimal if one cannot increase one party's profit without decreasing the other one's. Following Fan and Sullivan (2018), a Pareto optimal equilibrium of assortment vector a^* and wholesale price vector w^* of the repeated game can be represented as the solution to the following maximization problem with incentive compatibility constraints:

$$\max_{a_{st}, w_{st}, s \in \mathcal{S}, c \in \mathcal{C}} \sum_{s,c} \omega_r \pi_r(a_{st}, w_{st}, z_{st}) + \omega_m \pi_m(a_{st}, w_{st}, z_{st}) \quad (3)$$

$$\begin{aligned} \text{s.t. } & \sum_{s,c} \left[\pi_r(a_{st}, w_{st}, z_{st}) + \sum_{\tau=1}^{\infty} (\phi \delta)^\tau E\{\pi_r(a^*(z_{st+\tau}), w^*(z_{st+\tau}), z_{st+\tau}) | z_{st}\} \right. \\ & \left. + \sum_{\tau_0=1}^{\infty} \sum_{\tau=\tau_0}^{\infty} \phi^{\tau-1} (1-\phi) \delta^\tau E\{\pi_r(a^{NE}(z_{st+\tau}), w^{NE}(z_{st+\tau}), z_{st+\tau}) | z_{st}\} \right] \\ & \geq \sum_{s,c} \left[\pi_r^d(w_{st}, z_{st}) + \sum_{\tau=1}^{\infty} \delta^\tau E\{\pi_r(a^{NE}(z_{st+\tau}), w^{NE}(z_{st+\tau}), z_{st+\tau}) | z_{st}\} \right], \quad (4) \end{aligned}$$

$$\begin{aligned} & \sum_{s,c} \left[\pi_m(a_{st}, w_{st}, z_{st}) + \sum_{\tau=1}^{\infty} (\phi \delta)^\tau E\{\pi_m(a^*(z_{st+\tau}), w^*(z_{st+\tau}), z_{st+\tau}) | z_{st}\} \right. \\ & \left. + \sum_{\tau_0=1}^{\infty} \sum_{\tau=\tau_0}^{\infty} \phi^{\tau-1} (1-\phi) \delta^\tau E\{\pi_m(a^{NE}(z_{st+\tau}), w^{NE}(z_{st+\tau}), z_{st+\tau}) | z_{st}\} \right] \\ & \geq \sum_{s,c} \left[\pi_m^d(a_{st}, z_{st}) + \sum_{\tau=1}^{\infty} \delta^\tau E\{\pi_m(a^{NE}(z_{st+\tau}), w^{NE}(z_{st+\tau}), z_{st+\tau}) | z_{st}\} \right], \quad (5) \end{aligned}$$

where ω_r and ω_m are the weights of the retailer's and the manufacturer's profits in the Pareto

optimal equilibrium outcome (up to normalization), a_{st} is assumed to be continuous and represents the probabilities the products are available in market s , and the superscripts d and NE denote the most profitable deviation and the Nash equilibrium strategies. z_{st} includes the demand and cost shifters and is assumed to follow a stationary first-order Markov process. Importantly, z_{st} cannot be influenced by the firms' actions a_{st} and w_{st} . Essentially, a Pareto optimal equilibrium of the repeated game can be written as the solution to a constrained optimization problem that maximizes a weighted sum of the players' profits such that each player has no incentive to deviate from the equilibrium. The solution to the optimization problem is Pareto optimal because one cannot increase a player's profit without decreasing another one's. Following Bernheim and Whinston (1990) and Fan and Sullivan (2018), I assume deviations in one market s will be punished by reverting to static Nash equilibrium in all markets \mathcal{S} . Furthermore, in this multi-category setting, deviations in one category c will be punished by reversion in all categories \mathcal{C} (here, cider and lager beer). The category subscript c is omitted from all terms for expositional simplicity.

Because the sequential timeline here differs from the simultaneous one in Fan and Sullivan (2018), I assume any deviation from the equilibrium will not be detected by the partner until the next period. That is, if a retailer chooses off-equilibrium assortments in this period, the manufacturer will still charge equilibrium wholesale prices in this period, and the punishment only begins in the next period. In practice, this could happen when a manufacturer does not immediately notice the assortment changes on the retailer's shelf and still charges the same wholesale prices. Thus, although the timeline differs here, the incentive constraints and equilibrium conditions in Fan and Sullivan (2018) apply. On the other hand, if the manufacturer does not offer a discount, the retailer still sells the manufacturer's cider in this period because the retailer moves first.

Following Fan and Sullivan (2018), I obtain the first-order conditions of the above constrained optimization problem (3)-(5):

$$\omega_r \frac{\partial \pi_{rst}}{\partial a_{st}} + \omega_m \frac{\partial \pi_{mst}}{\partial a_{st}} + \gamma_{rt} \frac{\partial \pi_{rst}}{\partial a_{st}} + \gamma_{mt} \left(\frac{\partial \pi_{mst}}{\partial a_{st}} - \frac{\partial \pi_{mst}^d}{\partial a_{st}} \right) = 0 \quad (6)$$

$$\omega_r \frac{\partial \pi_{rst}}{\partial w_{st}} + \omega_m \frac{\partial \pi_{mst}}{\partial w_{st}} + \gamma_{rt} \left(\frac{\partial \pi_{rst}}{\partial w_{st}} - \frac{\partial \pi_{rst}^d}{\partial w_{st}} \right) + \gamma_{mt} \frac{\partial \pi_{mst}}{\partial w_{st}} = 0 \quad (7)$$

where $\gamma_{rt} \geq 0$ and $\gamma_{mt} \geq 0$ are the Karush-Kuhn-Tucker (KKT) multipliers of the retailer's and the manufacturer's incentive compatibility constraints. The discount factor δ , the relationship continua-

tion probability ϕ , and the profit functions π^* and π^{NE} enter into the first-order conditions through γ . When δ and ϕ are larger, γ would be smaller, and a higher coordination level can be sustained. Specifically, $\gamma = 0$ when the corresponding constraint is not binding, and both $\gamma = \infty$ at the static Nash equilibrium, because $\frac{\partial \pi_{rst}}{\partial a_{st}} = 0$, $\frac{\partial \pi_{mst}}{\partial a_{st}} = \frac{\partial \pi_{mst}^d}{\partial a_{st}}$, and $\gamma_{rt} = \left(\omega_r \frac{\partial \pi_{rst}}{\partial a_{st}} + \omega_m \frac{\partial \pi_{mst}}{\partial a_{st}} \right) / \frac{\partial \pi_{rst}}{\partial a_{st}} = \infty$ for equation (6), and similarly for equation (7). The multipliers do not have the market subscript s because the constraints are pooled across markets.

The above conditions can be rewritten as follows:

$$\frac{\partial \pi_{rst}}{\partial a_{st}} + \theta_{rmt} \frac{\partial \pi_{mst}}{\partial a_{st}} - \rho_{mt} \theta_{rmt} \frac{\partial \pi_{mst}^d}{\partial a_{st}} = 0 \quad (8)$$

$$\frac{\partial \pi_{mst}}{\partial w_{st}} + \theta_{mrt} \frac{\partial \pi_{rst}}{\partial w_{st}} - \rho_{rt} \theta_{mrt} \frac{\partial \pi_{rst}^d}{\partial w_{st}} = 0 \quad (9)$$

where $\theta_{rmt} = \frac{\omega_m + \gamma_{mt}}{\omega_r + \gamma_{rt}}$, $\theta_{mrt} = 1/\theta_{rmt}$, $\rho_{mt} = \frac{\gamma_{mt}}{\omega_m + \gamma_{mt}}$, and $\rho_{rt} = \frac{\gamma_{rt}}{\omega_r + \gamma_{rt}}$. These conditions are equivalent to the first-order conditions of the following static profit maximization problems, which are a tractable representation of the repeated game:

$$\max_{a_{st}} \pi_r(a_{st}, w_{st}, z_{st}) + \theta_{rmt} \pi_m(a_{st}, w_{st}, z_{st}) - \rho_{mt} \theta_{rmt} \pi_m^d(a_{st}, z_{st}) \quad (10)$$

$$\max_{w_{st}} \pi_m(a_{st}, w_{st}, z_{st}) + \theta_{mrt} \pi_r(a_{st}, w_{st}, z_{st}) - \rho_{rt} \theta_{mrt} \pi_r^d(w_{st}, z_{st}) \quad (11)$$

As discussed in Fan and Sullivan (2018), the parameters ρ and θ can be interpreted as a pair's coordination and asymmetry levels. $\rho \in [0, 1]$ is an increasing function of γ and hence inversely related to the discount factor δ , the relationship continuation probability ϕ , and the pair's coordination level. Intuitively, when a firm is not too worried about its partner's deviation (small ρ), due to high trust or a high discount factor, their coordination level would be high. Meanwhile, θ reflects the negotiation power of a firm relative to its partner, captured by its position in the equilibrium (measured by its weight ω in (3)) and the gain from relaxing its incentive constraint (measured by its KKT multiplier γ) relative to its partner's. Intuitively, when a firm's position is bad (small ω) or its incentive constraint is not binding ($\gamma = 0$), the firm has less negotiation power relative to its partner. In the special case where $\theta = 1$ and both $\rho = 0$, the solution corresponds to the vertical integration outcome. When both $\rho = 1$, which implies $\gamma = \infty$, the solution approaches the static Nash equilibrium.

Compared to the common profit-weight approach, this micro-founded model has an additional term that takes into account how a firm’s action would affect its partner’s deviation profit. Simulation results in Fan and Sullivan (2018) show that the profit-weight approach approximates but is not generally equivalent to this model—they are equivalent only if there is no coordination or perfect coordination. Although in principle time-specific θ_t and ρ_t can be estimated with data from each period, I assume time-invariant θ and ρ due to the persistence of beer shares. This modeling choice assumes away any equilibrium migration during the sample period, although the games are likely to have multiple equilibria.

To extend from one manufacturer-retailer pair to multiple pairs, I assume each pair holds all other pairs’ relationships fixed. In line with the bilateral bargaining literature, I assume passive beliefs, i.e., that a firm would believe other pairs still follow their equilibrium actions upon seeing a deviation. In the rest of the section, I present the empirical model in reverse order (demand, retail pricing, wholesale pricing, and assortment choice).

4.2 Demand

Demand is characterized by a random coefficient logit model (Berry et al., 1995). Consumer i derives the following utility from the purchase of cider j in store s and month t :

$$u_{ijst} = x_{jst}^1 \beta_i + \alpha_i p_{jst} + x_{jst}^2 \zeta + \xi_{jst} + \epsilon_{ijst}, \quad (12)$$

where p_{jst} is the price, x_{jst}^1 is a vector of product characteristics, x_{jst}^2 is a vector of marketing mix variables, ξ_{jst} is the product-store-month specific shocks, and ϵ_{ijst} is a type 1 extreme value error. The random coefficients (α_i, β_i) follow a multivariate normal distribution with mean $(\bar{\alpha}, \bar{\beta})$ and variance $\Sigma = \begin{pmatrix} \sigma_\alpha^2 & 0 \\ 0 & \sigma_\beta^2 \end{pmatrix}$. The vector x_{jst}^1 consists of three cider attributes: *abv*, alcohol by volume, *sugar*, amount of sugar per serving, and *seasonal*, whether the product is seasonal. These attributes could vary across time and stores because the data is aggregated to the brand-size level. The vector x_{jst}^2 consists of four variables: F_{jst} , the dummy for feature, D_{jst} , the dummy for display, $\log(Ads_{jd(s)t})$, TV ads in DMA d measured by gross rating points, and $\log(Ads_{all_d(s)t})$, TV ads for all cider brands in DMA d . I include the marketing mix variables because in-store promotion and advertising are often used to support new products. Consumers in each store choose among the

cider products carried by the store and the outside option of no purchase.

To construct a realistic counterfactual demand and to capture the growth and seasonality of hard cider, I include four types of fixed effects. Specifically, $\xi_{jst} = \xi_{jt} + \xi_{b(j)r(s)} + \xi_{b(j)d(s)} + \xi_s + \Delta\xi_{jst}$. The product-month fixed effects ξ_{jt} absorb the category expansion, seasonality, and each brand's entry time and growth path. The brand-retailer fixed effects ξ_{br} absorb the different brand preferences across retailers. The brand-DMA fixed effects ξ_{bd} capture the local taste and the local advertising intensity. The store fixed effects ξ_s capture the store-specific preferences for the whole category, which shape the substitution towards the outside option. These fixed effects capture the majority of variations in the data and are important for the identification of the parameters.

The share of cider j in store s and month t is given as follows:

$$s_{jst} = \int \frac{\exp(x_{jst}^1 \beta_i + \alpha_i p_{jst} + x_{jst}^2 \zeta + \xi_{jst})}{1 + \sum_k \exp(x_{kst}^1 \beta_i + \alpha_i p_{kst} + x_{kst}^2 \zeta + \xi_{kst})} dF(\alpha_i, \beta_i) \quad (13)$$

4.3 Supply

4.3.1 Retail Pricing

In the retail pricing stage, retailers set monopolist retail prices to maximize each store's profits:

$$\max_{p_{jst}, j \in \mathcal{J}_{st}} \pi_{rst} = \sum_{j \in \mathcal{J}_{st}} (p_{jst} - w_{jst}) s_{jst} N_s, \quad (14)$$

where \mathcal{J}_{st} is the set of products carried by store s in month t , w_{jst} is the wholesale price charged by the manufacturer, and N_s is the market size. I assume the retailers are monopolists and they know the demand shocks $\Delta\xi_{jst}$ when setting retail prices. w_{jst} reflects the effective wholesale price, which includes the price of the good and the transportation cost from the local warehouse to the store. Thus, w_{jst} can differ across stores, even in states that require wholesalers to sell at non-discriminatory prices.

Based on the descriptive evidence of no correlation between past beer shares and cider retail prices in the regulated states, I assume retailers maximize their static profits and do not take into account their relationships with manufacturers when setting retail prices. Further, negotiation over retail prices between manufacturers and retailers was considered per se illegal before Leegin (2007) and still involves a lot of uncertainty today, especially for firms with large market power (Steiner,

2010; Gundlach and Krotz, 2020).

Using the first-order condition of the retail profit function (14), I can invert out the wholesale price vector w_{st} :

$$w_{st} = p_{st} + [\Delta_{st}^r]^{-1} s_{st}, \quad (15)$$

where Δ_{st}^r is a $|\mathcal{J}_{st}|$ by $|\mathcal{J}_{st}|$ matrix with the (j, j') term equal to $\frac{\partial s_{j'st}}{\partial p_{jst}}$, and p_{st} and s_{st} are vectors of retail prices and market shares.

4.3.2 Wholesale Pricing

In the wholesale pricing stage, manufacturers set wholesale prices, taking into account their relationships with retailers and institutional constraints. In the regulated states, they engage in a static Nash-Bertrand game and do not offer preferential discounts. In the unregulated states, they take into account their relationships with retailers and set prices according to the objective function (11). For each store s and month t , they solve the following maximization problem:

$$\begin{aligned} & \max_{w_{jst}, j \in \mathcal{J}_{mst}} \pi_{mst} + \theta_{mr}^w d_s \pi_{rst} - \rho_{r,m}^w \theta_{mr}^w d_s \pi_{rst}^d \\ & = \sum_{j \in \mathcal{J}_{mst}} (w_{jst} - mc_{jst}) s_{jst} N_s + \theta_{mr}^w d_s \sum_{j \in \mathcal{J}_{st}} (p_{jst} - w_{jst}) s_{jst} N_s - \rho_{r,m}^w \theta_{mr}^w d_s \sum_{j \in \mathcal{J}_{st}^d} (p_{jst}^d - w_{jst}^d) s_{jst}^d N_s, \end{aligned} \quad (16)$$

where $\theta_{mr}^w = \frac{\omega_r^{mr} + \gamma_r^{mr}}{\omega_m^{mr} + \gamma_m^{mr}}$, $\rho_{r,m}^w = \frac{\gamma_r^{mr}}{\omega_m^{mr} + \gamma_m^{mr}}$, and the superscript mr denotes the game between manufacturer m and retailer r . \mathcal{J}_{mst} is the set of products offered by m in store s and month t , mc_{jst} is the marginal cost of j , d_s is a dummy variable that equals one if store s is in an unregulated state, and π_{rst}^d is r 's profit if r deviates. The model assumes that manufacturers know about their own and competitors' demand shocks $\Delta \xi_{jst}$ and take into account how retail prices would be set in the following stage when setting wholesale prices. If r deviates in the assortment stage, because punishment would not take place until the next period, the manufacturer would charge the equilibrium prices for products currently on the shelf (i.e., for $j \in \mathcal{J}_{st}$). For products currently not on the shelf (i.e., for $j \in \mathcal{J}_{st}^d \setminus \mathcal{J}_{st}$), the manufacturer would charge the static profit-maximizing prices.

The marginal costs mc_{jst} consist of production costs, transportation costs, store-specific supply costs, and cost shocks. I use product-month fixed effects λ_{jt} to capture the production costs, which depend on apple prices and could fluctuate over time, and I use product-county fixed effects λ_{jc} to

capture the transportation costs, which depend on manufacturer-to-store distances. I also use store fixed effects λ_s to capture the costs of selling at each store.

$$mc_{jst} = \lambda_{jt} + \lambda_{jc(s)} + \lambda_s + \eta_{jst}. \quad (17)$$

Using the first-order condition of the manufacturer's objective function (16), I can invert out the marginal cost vector mc_{st} as a function of θ_{mr}^w and $\rho_{r,m}^w$:

$$mc_{st} = w_{st} + [\Omega_{st}\Delta_{st}^w]^{-1} \left[s_{st} - \theta_{mr}^w d_s s_{st} + \rho_{r,m}^w \theta_{mr}^w d_s s_{st}^d \right], \quad (18)$$

where Δ_{st}^w is a $|\mathcal{J}_{st}|$ by $|\mathcal{J}_{st}|$ matrix with the (j, j') term equal to $\frac{\partial s_{j'st}}{\partial w_{jst}} = \sum_{j'' \in \mathcal{J}_{st}} \frac{\partial s_{j'st}}{\partial p_{j''st}} \frac{\partial p_{j''st}}{\partial w_{jst}}$, Ω_{st} is a $|\mathcal{J}_{st}|$ by $|\mathcal{J}_{st}|$ ownership matrix with the (j, j') term equal to 1 if j and j' are produced by the same manufacturer and zero otherwise.¹¹ The second term on the right is the negative of the wholesale markup.

From equation (18), we can see the wholesale markup mostly increases in $\rho_{r,m}^w$ and decreases in θ_{mr}^w since the term $[\Omega_{st}\Delta_{st}^w]^{-1}$ is negative. As discussed in Section 4.1, these parameters capture the coordination level (inversely) and the negotiation power of r relative to m . Intuitively, a low $\rho_{r,m}^w$ implies high coordination, and hence a low wholesale markup and less double marginalization. Similarly, a high θ_{mr}^w implies that r has large negotiation power relative to m , and hence a low wholesale markup and less double marginalization. When $\rho_{r,m}^w = 1$, the repeated game between m and r approaches the static Nash equilibrium, and thus $s_{st}^d = s_{st}$, and equation (18) reduces to the standard Nash-Bertrand pricing. When $\rho_{r,m}^w = 0$, r 's deviation profit does not matter, and the wholesale markup depends on θ_{mr}^w only. In this case, θ_{mr}^w should be less than 1 for the wholesale markup to be positive. Finally, for products currently on the shelf but that would be dropped had the retailer deviated (i.e., $s_{jst}^d = 0$), θ_{mr}^w also reflects the percentage discounts.

Combining equations (17) and (18) gives the following equation to be estimated:

$$w_{jst} + \Delta_{jst}^w = \theta_{mr}^w d_s \Delta_{jst}^w - \rho_{r,m}^w \theta_{mr}^w d_s \Delta_{jst}^{w,d} + \lambda_{jt} + \lambda_{jc} + \lambda_s + \eta_{jst}, \quad (19)$$

¹¹The derivative of retailer profit π_{rst} with respect to wholesale price w_{st} is $-s_{st} + \Delta_{st}^w(p_{st} - w_{st}) + \Delta_{st}^p s_{st} = -s_{st} + \Delta_{st}^p \Delta_{st}^r (-[\Delta_{st}^r]^{-1} s_{st}) + \Delta_{st}^p s_{st} = -s_{st}$, where $\Delta_{st}^w = \Delta_{st}^p \Delta_{st}^r$ and Δ_{st}^p is a $|\mathcal{J}_{st}|$ by $|\mathcal{J}_{st}|$ matrix with the (j, j') term equal to $\frac{\partial p_{j'st}}{\partial w_{jst}}$. Similarly, the derivative of π_{rst} to w_{st} is $-s_{st}^d$.

where $\Delta_{jst}^w = \left([\Omega_{st} \Delta_{st}^w]^{-1} s_{st} \right)_j$ is the negative of static Nash-Bertrand wholesale markup and $\Delta_{jst}^{w,d} = \left([\Omega_{st} \Delta_{st}^w]^{-1} s_{st}^d \right)_j$. Based on the descriptive findings, I assume $\rho_{r,m}^w$ is a function of past beer shares. Specifically:

$$\theta_{mr}^w = \theta^w; \quad \rho_{r,m}^w \theta^w = (\rho_0^w + \rho_1^w \times beer_share_{mr}^{2010}) \theta^w. \quad (20)$$

ρ_1^w is expected to be negative because the coordination level should be positively related to the past beer shares. Craft cider manufacturers' ρ^w and θ^w are assumed to be zero.

4.3.3 Assortment Choices

In the assortment stage, retailers choose the assortments for each store with the following objective:

$$\max_{\mathcal{J}_{st} \subseteq \mathcal{J}_t} \pi_r(\mathcal{J}_{st}) + \sum_m \theta_{rm}^a \pi_m(\mathcal{J}_{st}) - \sum_m \rho_{m,r}^a \theta_{rm}^a \pi_m^d(\mathcal{J}_{st}) - R_{st}, \quad (21)$$

where $\theta_{rm}^a = \frac{\omega_m^{mr} + \gamma_m^{mr}}{\omega_r^{mr} + \gamma_r^{mr}}$, $\rho_{m,r}^a = \frac{\gamma_m^{mr}}{\omega_m^{mr} + \gamma_m^{mr}}$, and the superscript mr denotes the game between manufacturer m and retailer r . R_{st} is the shadow cost to the retailer for stocking a cider assortment of size $|\mathcal{J}_{st}|$ in store s and month t , and \mathcal{J}_t is the set of products available in month t . Note that the objective function (21) pools retailer r 's objective (10) in each relationship. The choice over assortments is assumed to be made before the demand shocks $\Delta \xi_{jst}$ and the cost shocks η_{jst} are realized. I assume $\Delta \xi_{jst}$ and η_{jst} take the expected values, instead of averaging the profits over possible values of $\Delta \xi_{jst}$ and η_{jst} , for tractability. This information assumption about demand and cost shocks and the fact that assortment is not a continuous decision are the main differences between the model for estimation and the one shown in Section 4.1.

Similar to the wholesale pricing stage, the probability that r carries m 's products mostly decreases with $\rho_{m,r}^a$ and increases with θ_{rm}^a . Intuitively, a low $\rho_{m,r}^a$ implies high coordination, and so r is more likely to carry m 's products. Similarly, a high θ_{rm}^a implies that m has large negotiation power relative to r , and so r is more likely to carry m 's products. When $\rho_{m,r}^a = 1$, the outcome approaches the static Nash equilibrium, and thus $\pi_m^d = \pi_m$, and r 's assortment choice is only based on its static profit π_r . When $\rho_{m,r}^a = 0$, m 's deviation profit does not matter, and the assortment choice depends on θ_{rm}^a only. Because shelf space is limited, a good relationship between m and r would have a negative externality on the product availability of m 's competitors.

I use a revealed preference approach for estimation and assume that the observed assortment maximizes the retailer’s objective function. To construct the retailer’s assortment choice set, I consider all one-step deviations that replace a product in the current set with a same-size alternative. Let \mathcal{A}_{st} denote the set of assortments that include the observed one and all such deviations. The observed assortment \mathcal{J}_{st} satisfies:

$$\mathcal{J}_{st} = \arg \max_{\mathcal{J}'_{st} \in \mathcal{A}_{st}} \left[\pi_r(\mathcal{J}'_{st}) + \sum_m \theta_{rm}^a \pi_m(\mathcal{J}'_{st}) - \sum_m \rho_{m,r}^a \theta_{rm}^a \pi_m^d(\mathcal{J}'_{st}) \right]. \quad (22)$$

In other words, \mathcal{J}_{st} is the optimal assortment, and no one-step replacement deviation yields a higher objective value than \mathcal{J}_{st} . The shadow shelf cost R_{st} is dropped in equation (22) because the assortment size remains the same in the deviations.¹²

Based on the descriptive findings, I assume $\rho_{m,r}^a$ and θ_{rm}^a are functions of past beer shares:

$$\theta_{rm}^a = 1/(\theta_0^a + \theta_1^a \times beer_share_{mr}^{2010}); \quad \rho_{m,r}^a = \rho_0^a + \rho_1^a \times beer_share_{mr}^{2010}. \quad (23)$$

Specifically, θ_{rm}^a is parameterized in a way that is consistent with $\theta^a \theta^w = 1$. ρ_1^a is expected to be negative because the coordination level should be positively related to the past beer shares. Craft cider manufacturers’ ρ^a and θ^a are assumed to be zero.

5 Estimation and Identification

5.1 Demand

The demand estimation follows Berry et al. (1995) and is at the product-store-month level. Specifically, a product is defined as a brand-size combination. The market size is defined as the monthly average of all beer and cider sales in a store, which should cover all potential cider sales in that store. In the structural model estimation, I include the top 42 products that account for 90% of the sales, and include stores that have feature and display information and are large enough for precise share computation (i.e., average quantity sold ≥ 10 per product per month). The final sample consists of

¹²While R mostly reflects the option value of carrying another product, R could also contain the shelf maintenance costs, which are potentially different across manufacturers. The difference should be small across cider manufacturers that are also brewers because the marginal costs of restocking cider should be similar across them. That said, if the shelf maintenance costs are indeed smaller for the largest brewers (i.e., ABI and MillerCoors) than for other brewers that also produce cider (i.e., Heineken and Boston Beer), then the relationship effects might be over-estimated.

1,410 stores and 433,210 observations. All volumes and prices are normalized to the 6-pack (72oz) equivalent. The key parameters to estimate are $\{\alpha, \beta, \zeta, \Sigma\}$.

The identification follows the standard identification argument for a random coefficient logit model estimated with aggregate demand data. To handle potential price endogeneity caused by the firm knowledge of demand shocks $\Delta\xi_{jst}$, I use the average price of the same product across other retailers, in other counties of the same state, as the price instrument. The rationale is that the instrument and the price variable share the same local costs, such as the state excise tax and local transportation costs. The fixed effects capture the majority of variations in the data, and the marketing mix variables are assumed to be exogenous conditional on the fixed effects. Specifically, the advertising effects are identified by the variations within brand-DMA across time, and the feature and display effects are identified by the variations within brand-retailer across stores and time.

To estimate random coefficients, I use instruments that capture how much competition a product faces in a market, namely, the counts of products with similar characteristics. The characteristics include 1) whether the product has the same pack size (as the focal product), e.g., 6-pack, 2) whether the average price difference across all periods is within fifty cents (from the focal product), on top of the same pack size, 3) whether the average price difference across all periods is within one dollar, on top of the same pack size, 4) whether the ABV difference is within 0.5%, on top of the same pack size, 5) whether the sugar difference is within 2 grams, on top of the same pack size, and 6) whether the product is seasonal. These instruments are likely to be exogenous because, given the rich fixed effects, the remaining demand shocks $\Delta\xi_{jst}$ are unlikely to be realized before the assortments are chosen. The changes in market shares as these instruments change identify the random coefficients.

5.2 Supply

The supply model is estimated sequentially in reverse order, starting from the retail pricing stage. I first back out the wholesale prices using the first-order condition of the retail profit function (15). Then I estimate the wholesale pricing parameters θ_{mr}^w and $\rho_{r,m}^w$ in equation (19), using a linear instrumental variable approach. Lastly, I estimate the assortment parameters θ_{mr}^a and $\rho_{m,r}^a$ in equation (22) via maximum likelihood. Throughout the supply estimation, I take the demand estimates and the results in previous steps as given.

Wholesale pricing. I use an instrumental variable approach to estimate θ^w , $\rho_0^w \theta^w$, and $\rho_1^w \theta^w$ in

equation (19). The terms Δ_{jst}^w and $\Delta_{jst}^{w,d}$ are endogenous because both the derivative matrix Δ_{st}^w and the market shares s_{st} and s_{st}^d depend on prices and the cost shocks η_{jst} . To solve this endogeneity problem, I use demand shifters (Berry and Haile, 2014) and past beer shares as instruments. Based on the suggestions in Fan and Sullivan (2018), I use the market size, unobserved characteristics $\Delta\xi_{jst}$, and the number of seasonal products by rivals in the same market as demand shifters, all interacting with the unregulated dummy d_s . Intuitively, as demand changes, different wholesale conduct would imply different changes in manufacturers' and retailers' profits. Hence, the way these profits change with demand identifies the conduct parameters. The first stage p -values of the excluded instruments' F -statistics are 0.00 for all three variables. Standard errors are clustered at the product-store level.

To compute the deviation shares s_{st}^d , I first estimate an approximate model of the full model in (19), using the profit-weight approach (i.e., setting $\rho^w = 0$). Then I use the estimates to simulate deviation prices and shares and estimate the full model. While I cannot prove this algorithm is a contraction mapping, I repeat the second step with the new estimates and find very close results.

Assortment choices. I use a revealed preference approach with maximum likelihood to estimate θ_0^a , θ_1^a , ρ_0^a , and ρ_1^a . Denote $\hat{\pi}_r(\mathcal{J}_{st}) + \sum_m \theta_{rm}^a \hat{\pi}_m(\mathcal{J}_{st}) - \sum_m \rho_{m,r}^a \theta_{rm}^a \hat{\pi}_m^d(\mathcal{J}_{st})$ as r 's true objective value associated with assortment \mathcal{J}_{st} . Assume the observed assortment \mathcal{J}_{st} maximizes this objective value. Hence, for any alternative assortment $\mathcal{J}'_{st} \in \mathcal{A}_{st}$,

$$\left[\hat{\pi}_r(\mathcal{J}_{st}) + \sum_m (\theta_{rm}^a \hat{\pi}_m(\mathcal{J}_{st}) - \rho_{m,r}^a \theta_{rm}^a \hat{\pi}_m^d(\mathcal{J}_{st})) \right] - \left[\hat{\pi}_r(\mathcal{J}'_{st}) + \sum_m (\theta_{rm}^a \hat{\pi}_m(\mathcal{J}'_{st}) - \rho_{m,r}^a \theta_{rm}^a \hat{\pi}_m^d(\mathcal{J}'_{st})) \right] \geq 0. \quad (24)$$

As described in Section 4.3.3, \mathcal{A}_{st} consists of the observed assortment and all one-step deviations that replace a product in the current set with a same-size alternative. Thus, the shelf space required is the same for all assortment choices in \mathcal{A}_{st} . Further, I restrict the replacement set to products available in the same county (i.e., sold in other stores within the county), to minimize the concern that the replacement product is not available in the region at all, instead of just not being selected by the retailer.¹³ In total, I construct 591,170 one-step replacement deviations from the 46,830 observed assortments at the store-month level. Note that, in principle, one could expand \mathcal{A}_{st} to include multi-step deviations and adding or removing a product. The latter would lead to changes in shelf space for the category and involve comparisons with products outside the category.

¹³To sell its hard cider in a state, a manufacturer needs to contract with a licensed wholesaler in the state. The manufacturer may find it better not to sell in the state/region at all if it cannot reach the economies of scale.

Following Crawford et al. (2018), I assume $\nu(\mathcal{J}_{st})$ to be the measurement error of the true objective value associated with assortment \mathcal{J}_{st} . That is,

$$\nu(\mathcal{J}_{st}) = \left[\hat{\pi}_r(\mathcal{J}_{st}) + \sum_m (\theta_{rm}^a \hat{\pi}_m(\mathcal{J}_{st}) - \rho_{m,r}^a \theta_{rm}^a \hat{\pi}_m^d(\mathcal{J}_{st})) \right] - \left[\pi_r(\mathcal{J}_{st}) + \sum_m (\theta_{rm}^a \pi_m(\mathcal{J}_{st}) - \rho_{m,r}^a \theta_{rm}^a \pi_m^d(\mathcal{J}_{st})) \right], \quad (25)$$

where $\nu(\mathcal{J}_{st})$ is type 1 extreme value, independent and identically distributed across assortment-store-month with scale parameter σ^ν . Because the model is flexible in capturing retailer-specific and location-specific factors in both demand and costs, it is reasonable to assume that the measurement error ν is orthogonal to the profits. The probability that \mathcal{J}_{st} is the observed assortment is thus:

$$\Pr(\mathcal{J}_{st}) = \frac{\exp \left[(\pi_r(\mathcal{J}_{st}) + \sum_m (\theta_{rm}^a \pi_m(\mathcal{J}_{st}) - \rho_{m,r}^a \theta_{rm}^a \pi_m^d(\mathcal{J}_{st}))) / \sigma^\nu \right]}{\sum_{\mathcal{J}'_{st} \in \mathcal{A}_{st}} \exp \left[(\pi_r(\mathcal{J}'_{st}) + \sum_m (\theta_{rm}^a \pi_m(\mathcal{J}'_{st}) - \rho_{m,r}^a \theta_{rm}^a \pi_m^d(\mathcal{J}'_{st}))) / \sigma^\nu \right]}. \quad (26)$$

With the parameterization of θ_{rm}^a and $\rho_{m,r}^a$ in equation (23), the log-likelihood is then:

$$LL(\theta_0^a, \theta_1^a, \rho_0^a, \rho_1^a, \sigma^\nu) = \sum_{st} \log(\Pr(\mathcal{J}_{st})). \quad (27)$$

I normalize the weight on the retailer's profits to one and instead estimate the scale of the measurement error σ^ν .

To compute the profits π_{rst} , π_{mst} , and π_{mst}^d for each deviation \mathcal{J}'_{st} , I construct the marketing and cost variables for the replacement product j' and solve for the new equilibrium wholesale and retail prices for all products in \mathcal{J}'_{st} . Specifically, the advertising variables are directly observed at the brand-DMA-month level. Feature and display variables are extrapolated to be the same retailer-month average because a retailer often uses the same promotion strategy across stores. The various fixed effects follow the model estimates. To compute π_{mst}^d , I simulate the wholesale prices for the manufacturer that deviates, which maximizes its current period profit, and keep the other manufacturers' prices unchanged.

The identifying variations for $\rho_{m,r}^a$ and θ_{rm}^a come from how retail assortments vary with the retailer's and manufacturer's profits across markets. For each retailer r , the profits π_{rst} , π_{mst} , and π_{mst}^d differ across stores. These differences are driven by different consumer preferences, costs, and local product availability across stores. Along with r 's assortment choices across stores, they provide variations to pin down $\rho_{m,r}^a$ and θ_{rm}^a . Specifically, r 's assortment responses to a manufacturer's

deviating profits reflect their coordination level. With the parameterization of θ_{rm}^a and $\rho_{m,r}^a$ as functions of past beer shares, r 's assortment responses and their correlation with past beer shares identify $\{\theta_0^a, \rho_0^a\}$ and $\{\theta_1^a, \rho_1^a\}$. The extent to which the assortment choices appear as random identifies σ^ν . These identifying variations mainly come from across-market variations rather than across-time variations, because the assortments of each store are relatively stable over time.

6 Results

6.1 Demand Estimates

Table 6 reports the demand estimates. Column (1) reports the logit specification, then columns (2)-(3) report the main specification with random coefficients for price and alcohol content.¹⁴ In the main specification, the mean price coefficient is negative and significant (-0.632), indicating that consumers buy less cider as the price increases. The random coefficient for price is large and marginally significant (0.187), suggesting a sizable heterogeneity in consumer price sensitivity. These estimates translate to an average own price elasticity of -3.89 , which is slightly smaller (in magnitude) than the one for lager beer in the literature (Miller and Weinberg, 2017). This difference can probably be explained by the fact that hard cider is more premium than lager beer.

There is also significant heterogeneity in consumer preference for alcohol content. While the average consumer prefers low-alcohol cider (-0.807), some consumers instead favor high-alcohol cider ($\text{sd} = 0.558$). Because the latter group is more likely to make a cider purchase, without this random coefficient, one might reach the wrong conclusion that the average consumer prefers high-alcohol cider (0.227 , column (1)). Consumers also like seasonal (0.476) and less-sugar (-0.014) cider.

Among the marketing supports, I find a positive and significant effect of in-store display (0.427). The effect size is comparable to a \$1 price drop (average \$9.55). Effects of the other three marketing variables, however, are either statistically insignificant (brand ads and category ads) or negative and statistically significant (feature). Note that the effects of these variables might not be well identified, because product-month fixed effects absorb most of the variations in the ad variables (most ads are

¹⁴An alternative specification includes four random coefficients for price, alcohol content, sugar, and seasonal status. It yields nearly identical estimates, and the latter two random coefficients are statistically insignificant.

Table 6: Demand Estimates

	(1) logit	(2) random coef. logit mean	sd
Price	-0.391*** (0.013)	-0.632*** (0.179)	0.187* (0.098)
Abv	0.227*** (0.012)	-0.807** (0.361)	0.558*** (0.085)
Seasonal	0.426*** (0.013)	0.476*** (0.017)	
Sugar (g)	-0.013*** (0.001)	-0.014*** (0.002)	
log(Ads)	-0.009 (0.013)	-0.012 (0.018)	
log(Ads, all)	-0.003 (0.009)	-0.007 (0.011)	
Feature	-0.006 (0.014)	-0.051*** (0.017)	
Display	0.401*** (0.006)	0.427*** (0.012)	
Avg Own Price Elasticity	-3.73	-3.89	
Product-Month FE	Yes	Yes	
Brand-Retailer FE	Yes	Yes	
Brand-DMA FE	Yes	Yes	
Store FE	Yes	Yes	
Observations	432,210	432,210	

Notes: The observation unit is brand-size-store-month. The first stage has a t-statistic of 73. *p<0.1; **p<0.05; ***p<0.01

national ads), and brand-retailer fixed effects absorb most of the variations in the feature variable. Nevertheless, their effect magnitude is much smaller than that of in-store display, which I find to have an economically meaningful effect on cider sales.

6.2 Supply Estimates


Pricing. Table 7 reports the estimates for wholesale pricing, and Figure 5 plots how the coordination parameter ρ^w changes with the past beer share based on these estimates.¹⁵ We can see that ρ^w decreases with the past beer share, implying that relationship trust and coordination increase with the past beer share. Specifically, for brewers with high market shares (i.e., ABI and MillerCoors), ρ^w is around 0.8 and different from both 0 and 1, suggesting a moderate level of coordination

¹⁵Ohio is excluded from the supply estimation because of its minimum markup requirement for alcoholic beverages. Retailers not in the 2010 data are also excluded.

Table 7: Pricing Estimates

θ^w	0.67*** (0.05)
$\rho_0^w \times \theta^w$	0.66*** (0.03)
$\rho_1^w \times \theta^w$	-0.40*** (0.09)
Average retail margin	35.9%
Average wholesale margin	34.2%
Average wholesale price (6-pack)	\$6.26

Notes: Standard errors are clustered at the product-store level. *p<0.1; **p<0.05; ***p<0.01

Figure 5: ρ^w and past beer share

between these brewers and retailers. For brewers with relatively low market shares (i.e., Heineken and Boston Beer), ρ^w is close to 1, suggesting little coordination between these brewers and retailers. The asymmetry parameter θ^w is estimated to be 0.67, which is less than 1, implying the negotiation power is relatively larger for the manufacturers than for the retailers.

These relationship estimates imply that ABI and MillerCoors offer average wholesale discounts of 7.6% and 5.5% to retailers in unregulated states, consistent with the descriptive evidence in Section 3. These discounts compensate retailers for their preferential treatments in assortments, at the same time reducing double marginalization. The average recovered retail margin ($\frac{p-w}{p}$) is 35.9%, and the implied average wholesale price is \$6.26 per 6-pack. The average recovered wholesale margin ($\frac{w-mc}{w}$) is 34.2%, and the wholesale margins account for 21.7% of the retail prices. As a robustness test, I find less than 0.06% of the recovered marginal costs are below zero.

Assortment Choices. Table 8 reports the parameter estimates for assortment choices, and Figure 6 plots how the coordination parameter ρ^a varies with the past beer share based on these estimates. Similar to Figure 5, ρ^a decreases with the past beer share, which suggests relationship trust and coordination increase with the past beer share. The coordination is moderate between the high-share brewers (ABI and MillerCoors) and retailers (ρ^a around 0.8). By contrast, the coordination is minimal between the relatively low-share brewers (Heineken and Boston Beer) and retailers (ρ^a around 1). These results imply that retailers carry more ciders from ABI and MillerCoors

than static profit-maximizing retailers would carry.¹⁶

The θ^a estimates suggest an asymmetry between manufacturers and retailers, but the asymmetry does not depend on past beer shares ($\theta_1^a = -0.004$). Because $\theta_{rm}^a = 1/(\theta_0^a + \theta_1^a \times \text{beer_share}_{mr}^{2010}) > 1$, these estimates suggest the manufacturers have relatively larger negotiating power than the retailers (similar to the result in wholesale pricing). On the other hand, the estimates of ρ_1^a and θ_1^a suggest past beer shares affect only the coordination, not the asymmetry.

Summary. In sum, these ρ estimates suggest that relationship trust and coordination increase with the past beer share. The coordination is moderate between retailers and the leading brewers, ABI and MillerCoors, but minimal between retailers and Heineken and Boston Beer. I also find the manufacturers have relatively larger negotiating power than the retailers. Note that although $\theta_{rm}^a = 1/\theta_{mr}^w$ theoretically, I only find $\theta_{rm}^a > 1$ and $1/\theta_{mr}^w > 1$, and they are not exactly equal to each other. In Web Appendix C, I present the ratios of ω and γ and show how they change with the past beer share, using the estimates of ρ and θ . In Web Appendix D, I present the model estimates using the alternative profit-weight approach.

7 Counterfactual: Impact of Manufacturer-Retailer Relationships

Using the model estimates, I conduct two sets of counterfactual simulations to quantify the impact of manufacturer-retailer relationships in the hard cider market. In the first set of simulations, I evaluate the relationship's impact on various market outcomes, including product distribution, profits, and welfare. I compare the current market scenario with a counterfactual scenario in which manufacturers and retailers maximize their static profits. Because the retailers and leading brewers coordinate in assortments and wholesale prices, their relationships' impact depends on the extent of assortment distortion towards the leading manufacturers and the degree of prices reduction. This impact reflects the difference between a truly competitive environment and the actual environment in which manufacturers and retailers resort to relationships when regulations restrict fixed transfers and contract forms.

In the second set of simulations, I compute the amounts of wholesale markup discounts and

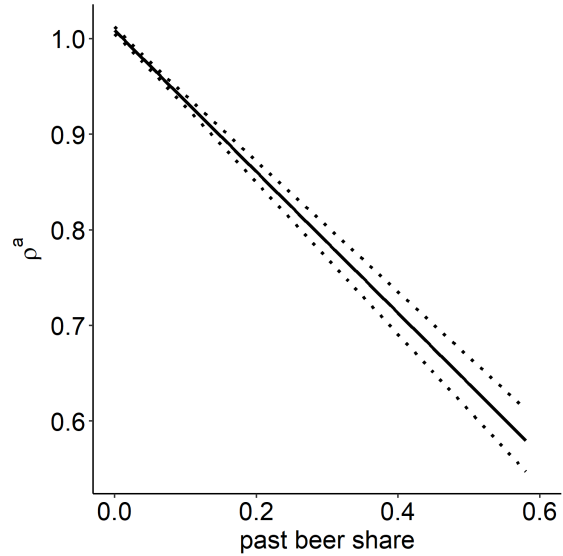
¹⁶A natural question is whether these results are driven by the ex-ante high expectations from retailers for ABI's and MillerCoors' products, but they end up not performing well. To answer this question, in a robustness check I restrict the sample to October 2014 onward, when all major brands have entered for at least six months and the demand has been stable. The results are close.

Table 8: Assortment Choices Estimates

θ_0^a	0.16*** (0.003)
θ_1^a	-0.004 (0.03)
ρ_0^a	1.01*** (0.002)
ρ_1^a	-0.74*** (0.03)
σ^ν	0.04*** (0.0004)

Notes: Standard errors in parentheses are based on the assortment stage only. *p<0.1; **p<0.05; ***p<0.01

Figure 6: ρ^a and past beer share



lump-sum transfers that are necessary to generate the same distribution effect of manufacturer-retailer relationships. The two counterfactuals highlight the importance of these relationships and the difference between the relationships and other vertical practices. Note that a limited set of counterfactuals can be evaluated because ρ and θ are a reduced-form characterization of the underlying repeated games.

7.1 Impact of Manufacturer-Retailer Relationships

I take the following steps to simulate market outcomes for the current scenario and the counterfactual scenario in which manufacturers and retailers maximize their static profits (i.e., $\theta = \rho = 0$). For each store, I first construct one-step replacement deviations from the current assortment to form the set of alternative assortments. For each deviation, the replacement product has to have the same size and be available in the store's county (i.e., sold by another store in the county). This criterion ensures that the replacement product is a readily available option for the retailer. Then, for each scenario and alternative assortment, including the current one, I solve for the equilibrium wholesale and retail prices using equations (15) and (18), along with the demand and supply estimates. I follow the same procedure as in Section 5.2 to construct the marketing and cost variables for the

replacement products. Because integrating profits over the demand and cost shock distributions is computationally costly, I set $\Delta\xi_{jst} = 0$ and $\eta_{jst} = 0$ for tractability, noting that the shocks explain small fractions of variations in the demand and supply models. Lastly, I pin down the assortments that would be chosen based on the retailers’ objectives in each scenario. Note that all manufacturers’ and retailers’ objectives will change in the counterfactual scenario, except for the manufacturers’ objectives in the regulated states.

Overall, manufacturer-retailer relationships provide a double win for both sides. Table 9 reports the first set of simulation results, presented as the percentage differences between the current scenario and the counterfactual scenario (baseline, no relationships). Specifically, ABI and MillerCoors gain 17.5% and 5.1% in distribution and 9.7% (\$0.5m) and 1% (\$0.1m) in profits from the relationships. Even in the presence of heavy regulations, the relationships can distort assortments and mitigate the impact of regulations. Retailers have a profit increase of 1.1% (\$2.1m) due to the gain from wholesale discounts, despite their loss from sub-optimal assortments. These sub-optimal assortments do not cost the retailers much, because marginal changes in assortments do not make much difference in their profits. However, these changes can have a big impact on manufacturers’ profits because the changes are “in or out” differences for them. As a whole, ABI, MillerCoors, and the retailers’ profits increase by 1.4% (\$2.7m). This bigger pie comes from the larger product availability for the two manufacturers and the lower double marginalization in the channel.

The gains of the largest brewers and retailers are at the expense of other cider manufacturers, which lose distribution and profits as a result of lacking relationships and being left out of coordination. Specifically, Heineken loses 7.9% (−\$0.6m) in profits, Boston Beer loses 0.9% (−\$0.6m), and craft cider manufacturers lose 2.7% (−\$0.9m). Together, they lose 4.1% in distribution and 1.9% (−\$2.1m) in profits. The current alcohol regulations fail to generate a fair playing field for them. Combining all manufacturers and retailers, the overall channel sees a slight profit increase of 0.2% (\$0.6m).

As for consumers, their surplus goes up by 1.1% (\$2.3m) because of the reduced double marginalization from the wholesale discounts, despite the biased assortments towards the leading manufacturers. This result, however, should be interpreted with caution in this context, due to the public health and social concerns inherent in alcohol consumption.

Relating these results to the literature provides additional insights into the contrast between

Table 9: Impact of Manufacturer-Retailer Relationships

		Relationship Impact (percentage change)	Absolute Change (\$m) (industry per year)
Anheuser-Busch InBev	Distribution	17.5%	
	Profit	9.7%	0.5
MillerCoors	Distribution	5.1%	
	Profit	1%	0.1
Heineken	Distribution	-7.7%	
	Profit	-7.9%	-0.6
Boston Beer	Distribution	-1.4%	
	Profit	-0.9%	-0.6
Craft Ciders	Distribution	-3.7%	
	Profit	-2.7%	-0.9
Retailers	Profit	1.1%	2.1
Retailers+ABI+MC	Profit	1.4%	2.7
HN+BB+Craft	Profit	-1.9%	-2.1
Channel (M+R)	Profit	0.2%	0.6
Consumer	Surplus	1.1%	2.3

Notes: The relationship impact is defined as the relative differences between the current scenario and the counterfactual scenario (baseline, no relationships). The distribution measure is averaged across products.

the relationships studied here and slotting fees. Hristakeva (2022b) finds that slotting fees increase retailers' profits by 3.3% and have little impact on consumer surplus relative to linear wholesale contracts. Here, the profit increase for retailers is smaller (1.1%) because they can only use pair relationships in this heavily regulated environment. However, the relationships can still recoup a sizable part of the retailers' missing profits. On the other hand, although assortments are worse in both Hristakeva (2022b) and the present paper, we find different effects on consumer surplus. This difference is driven by the direct effect on wholesale prices and double marginalization from the relationships. Thus, from the regulators' perspective, pair relationships are still better than slotting fees.

While these relationships are hard to regulate, policy makers can try to break them down by cutting the channels the pairs secretly exchange benefits. In this specific case, implementing posted wholesale prices across all states (and not just regulated states) could break them down by preventing manufacturers from compensating their preferred retailers. Because most of these retailers are multi-state retailers that operate in both types of states, they would stop biasing assortments towards the leading manufacturers in all states once they cannot get compensation from them in

the unregulated states. Thus, the relationships' impact presented in Table 9 will also inform the impact of a nationwide tightened wholesale pricing law, which is a combination of higher prices in the currently unregulated states and better assortments across all states.

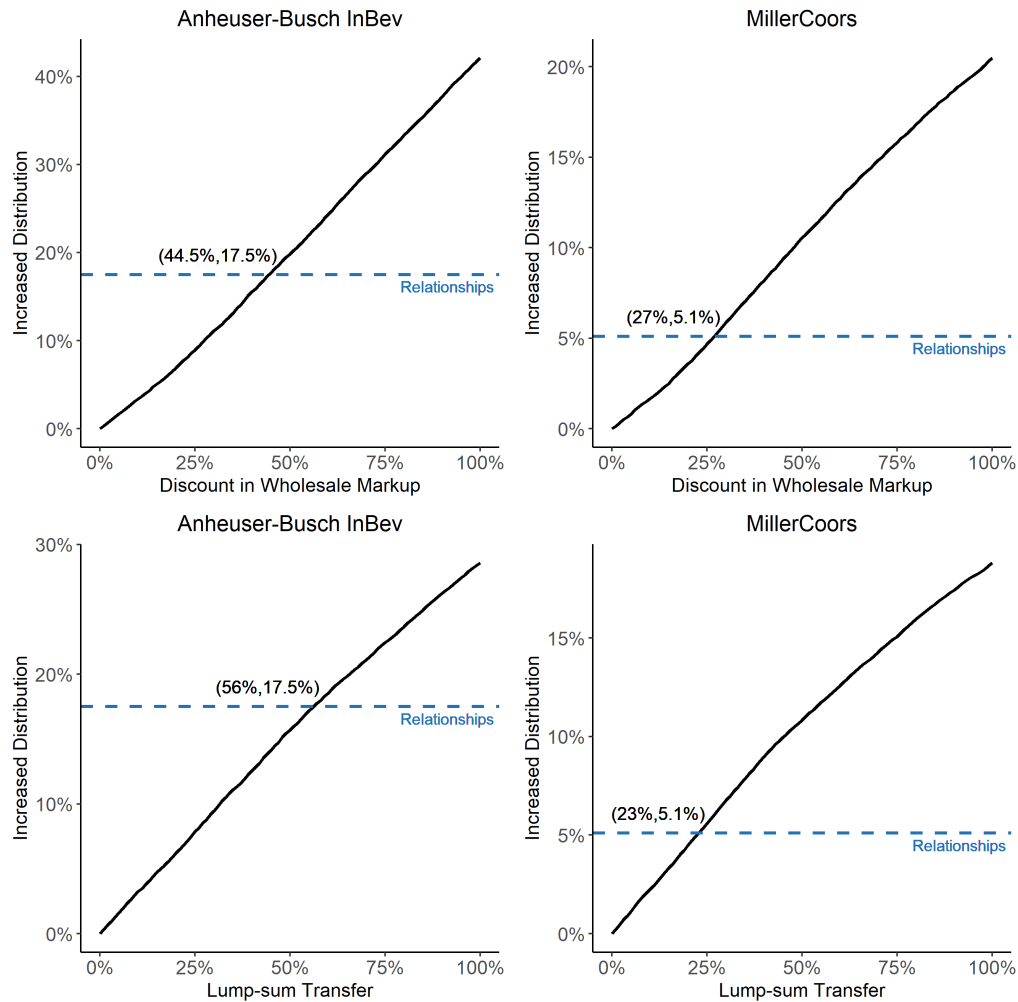
Finally, the generalizability of these findings to other settings depends on their market structures and regulations. Specifically, manufacturer-retailer relationships can be very important for markets with dominant manufacturers and high concentration, which is often the case for CPG markets. In addition, regulations determine the set of tools manufacturers and retailers can collaborate on. In general, the relationships' impacts estimated in this paper best apply to duopoly markets where wholesale discounts are the main subsidy offered by manufacturers.

7.2 Equivalent Wholesale Markup Discounts and Lump-sum Transfers

I conduct a second set of simulations to conceptualize the relationships' effects on product distribution. Specifically, I examine how large the wholesale markup discounts and lump-sum transfers need to be to secure the current product distribution for leading manufacturers if they do not have specific relationships with retailers. Both of these strategies have been widely used by manufacturers to increase product distribution. Although lump-sum transfers are illegal in the alcohol category, they are anecdotally common in others. On the other hand, manufacturers can offer non-discriminatory wholesale discounts in all states, separately from their specific relationships with retailers. Instead of simulating a market where manufacturers compete for shelf space with committed wholesale discounts or lump-sum transfers, for which an equilibrium might not exist, I focus on the amount of transfers and discounts needed for one manufacturer to generate the same distribution level, assuming the others do not make such an offer.

The steps to simulate these discounts and transfers are as follows. For the wholesale markup discounts, I first compute the percentage of stores that would carry the focal manufacturer's products as its discounts move from 0 to 100%. I assume the other manufacturers charge the same wholesale prices as in the baseline, no-relationship scenario, and retailers will re-optimize retail prices for the new wholesale prices. I then pin down the level of discounts that would generate the current distribution level for the focal manufacturer. For the lump-sum transfers the steps are similar, except that lump-sum transfers do not directly affect prices. I assume the transfers take the form of a fixed percentage of the focal manufacturer's profits. In practice, lump-sum transfers such as

Figure 7: Equivalent Wholesale Markup Discounts and Lump-sum Transfers



Notes: The focal manufacturer’s distribution increase is measured by the change in the percentage of stores that would carry its products, averaged across products. The distribution change is calculated for each 1% bin of discounts and transfers. Lump-sum transfers take the form of a fixed percentage of the manufacturer’s profits.

slotting fees and trade allowances can be either a fixed amount of payments or a fixed percentage of the manufacturer’s revenues. This simulation exercise is more in line with the latter format.

Figure 7 shows that only deep wholesale discounts and lump-sum transfers can generate the same distribution level as manufacturer-retailer relationships do. Wholesale markups would have to be 44.5% lower for ABI and 27% lower for MillerCoors across all stores, including the ones in the regulated states. These discounts are about twice as large as the ones they are currently offering in the unregulated states only. For lump-sum transfers, ABI would have to offer 56% of its wholesale profits, and MillerCoors would have to offer 23%. These numbers could grow even higher if we also

allow other manufacturers to offer transfers and discounts.

These results highlight the strategic importance of manufacturer-retailer relationships in promoting new products and maintaining market-dominant positions for leading CPG manufacturers. As many manufacturers face stagnation and decline in their main products, their competitive advantage from their close relationships with retailers can help them thrive in new product markets and categories. As a result, maintaining close relationships with retailers is a good strategy for them to extend their dominance, and this can contribute to the persistently high concentration in CPG markets.

8 Conclusion

In this paper, I present empirical evidence of manufacturer-retailer relationships and develop a repeated game-based model to estimate their impact in the unique setting of the US hard cider market. I show that retailers and leading manufacturers coordinate with and offer preferential treatments to each other when setting assortments and wholesale prices. These coordinations and relationships distort assortments towards leading manufacturers while reducing double marginalization. Although these relationships could improve welfare, the results imply that the current regulations in the alcoholic beverage industry do not successfully generate a level playing field for every manufacturer.

This paper can be extended in two ways. First, it focuses on preferential assortment choices and wholesale prices as two cooperation channels of the relationships. In future work, it would be interesting to explore other channels, such as information exchange. Also, it would be interesting to explore the equilibrium outcomes when manufacturers have access to these relationships as well as other types of vertical arrangements, such as rebates and fixed transfers. Second, this paper characterizes present, established manufacturer-retailer relationships. Understanding how these relationships and the firms' beliefs develop and change over time, via communication and relationship-specific investments, would be another important extension.

Funding and Competing Interests

The author certifies that the author has no affiliations with or involvement in any organization or entity with any financial interest or non-financial interest in the subject matter or materials discussed

in this manuscript. The author has no funding to report.

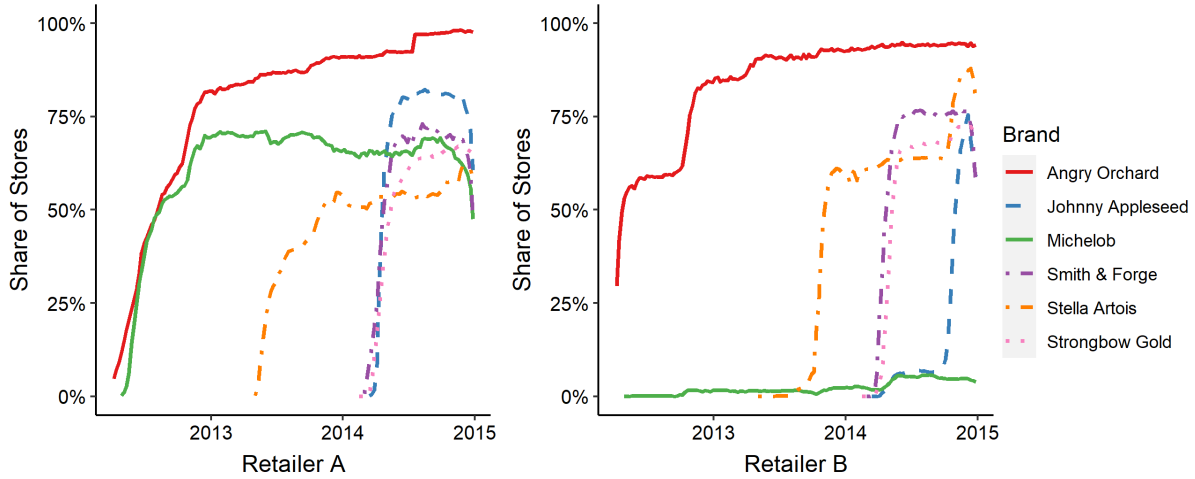
References

- Asker, J. (2016). Diagnosing foreclosure due to exclusive dealing. *The Journal of Industrial Economics* 64(3), 375–410.
- Bernheim, B. D. and M. D. Whinston (1990). Multimarket contact and collusive behavior. *The RAND Journal of Economics*, 1–26.
- Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society*, 841–890.
- Berry, S. T. and P. A. Haile (2014). Identification in differentiated products markets using market level data. *Econometrica* 82(5), 1749–1797.
- Bronnenberg, B. J. and C. F. Mela (2004). Market roll-out and retailer adoption for new brands. *Marketing Science* 23(4), 500–518.
- Conlon, C. T. and J. H. Mortimer (2021). Efficiency and foreclosure effects of vertical rebates: Empirical evidence. *Journal of Political Economy* 129(12), 3357–3404.
- Crawford, G. S., R. S. Lee, M. D. Whinston, and A. Yurukoglu (2018). The welfare effects of vertical integration in multichannel television markets. *Econometrica* 86(3), 891–954.
- Draganska, M., D. Klapper, and S. B. Villas-Boas (2010). A larger slice or a larger pie? an empirical investigation of bargaining power in the distribution channel. *Marketing Science* 29(1), 57–74.
- Dwyer, F. R., P. H. Schurr, and S. Oh (1987). Developing buyer-seller relationships. *Journal of marketing* 51(2), 11–27.
- Ellickson, P. B., P. Kong, and M. J. Lovett (2018). Private labels and retailer profitability: Bilateral bargaining in the grocery channel. *Available at SSRN 3045372*.
- Fan, Y. and C. Sullivan (2018). Measuring competition with a flexible supply model. Working Paper.
- Fan, Y. and C. Yang (2020). Competition, product proliferation and welfare: A study of the us smartphone market. *American Economic Journal: Microeconomics* 12(2), 99–134.
- Gross, A. (2019). Private labels and bargaining in the supply chain.
- Gundlach, G. T. and P. N. Bloom (1998). Slotting allowances and the retail sale of alcohol beverages. *Journal of Public Policy & Marketing* 17(2), 173–184.
- Gundlach, G. T. and R. T. Krotz (2020). Resale price maintenance: implications of marketing trends for the colgate doctrine and the leegin factors. *Journal of Public Policy & Marketing* 39(1), 48–61.
- Heide, J. B. (1994). Interorganizational governance in marketing channels. *Journal of marketing* 58(1), 71–85.
- Hristakeva, S. (2022a). Determinants of channel profitability: retailers’ control over product selections as contracting leverage. *Marketing science* 41(2), 315–335.

- Hristakeva, S. (2022b). Vertical contracts with endogenous product selection: An empirical analysis of vendor allowance contracts. *Journal of Political Economy* 130(12), 3202–3252.
- Hwang, M., B. J. Bronnenberg, and R. Thomadsen (2010). An empirical analysis of assortment similarities across us supermarkets. *Marketing Science* 29(5), 858–879.
- Kaufman, P., S. Jayachandran, and R. L. Rose (2006). The role of relational embeddedness in retail buyers’ selection of new products. *Journal of Marketing Research* 43(4), 580–587.
- Leegin Creative Leather Products, Inc. v. PSKS, Inc. (2007). 551 US 877.
- Miller, N. H. and M. C. Weinberg (2017). Understanding the price effects of the millercoors joint venture. *Econometrica* 85(6), 1763–1791.
- Misra, K. (2008). Understanding retail assortments in competitive markets.
- Steiner, R. L. (2010). The leegin factors-a mixed bag. *The Antitrust Bulletin* 55(1), 25–58.
- Sudhir, K. and V. R. Rao (2006). Do slotting allowances enhance efficiency or hinder competition? *Journal of Marketing Research* 43(2), 137–155.
- Villas-Boas, S. B. (2007). Vertical relationships between manufacturers and retailers: Inference with limited data. *The Review of Economic Studies* 74(2), 625–652.
- Viswanathan, M., O. Narasimhan, and G. John (2020). Economic impact of category captaincy: an examination of assortments and prices. *Marketing Science*.
- Weitz, B. A. and S. D. Jap (1995). Relationship marketing and distribution channels. *Journal of the academy of Marketing Science* 23(4), 305–320.
- Zhang, J. Z., O. Netzer, and A. Ansari (2014). Dynamic targeted pricing in b2b relationships. *Marketing Science* 33(3), 317–337.
- Zhang, J. Z., G. F. Watson Iv, R. W. Palmatier, and R. P. Dant (2016). Dynamic relationship marketing. *Journal of Marketing* 80(5), 53–75.

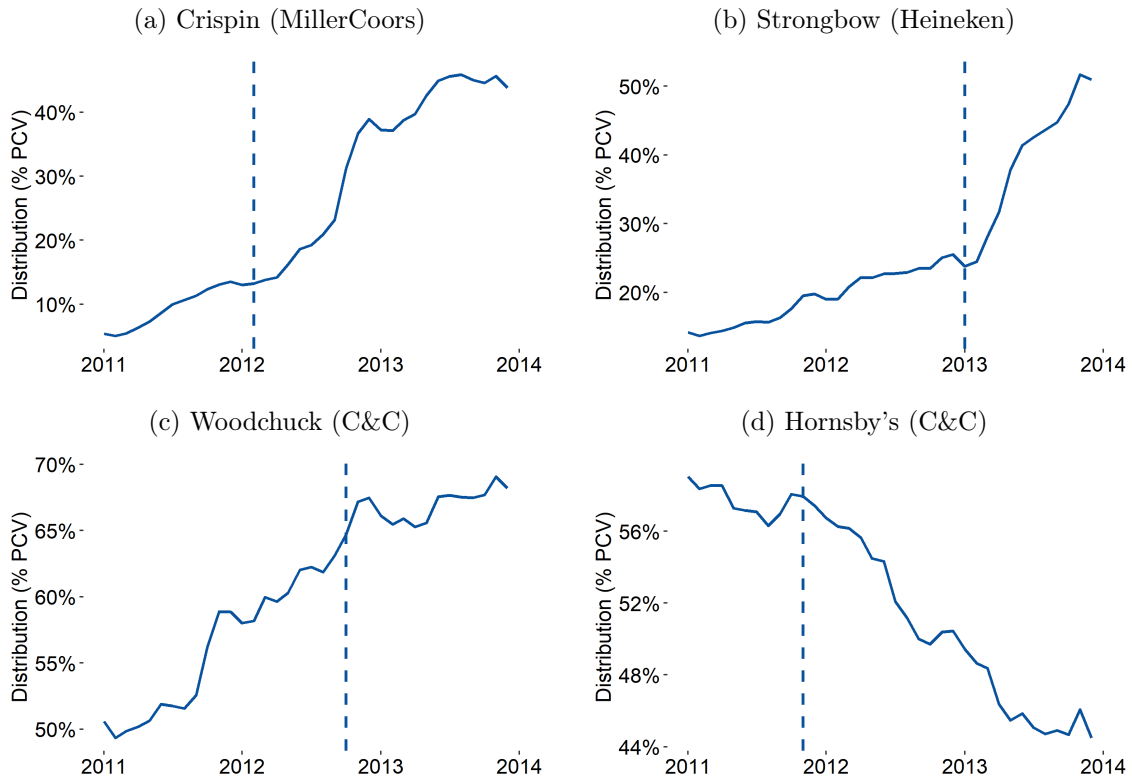
Appendix

Figure 8: Adoption of Major Brewers' New Ciders: Top Two Grocery Retailers



Notes: The figure shows the top two grocery retailers' cider adoption patterns in the sample. Both retailers quickly adopted major brewers' new ciders after their launches. The vertical part of the lines suggests that most stores moved around the same time for each retailer. By contrast, many brands' adoption levels and timing differ across retailers. For example, Michelob was well adopted by Retailer A but not by Retailer B, and Johnny Appleseed was immediately adopted by Retailer A but not by Retailer B. Such adoption differences across retailers could be driven by the differences in their relationships with the brewers.

Figure 9: Distribution of Incumbent Brands around Ownership Changes



Notes: Crispin (a) and Strongbow (b) saw a jump in distribution after their acquisitions by a major US brewer. Woodchuck (c) and Hornsby's (d) did not see such a jump after being acquired by the Irish cider company C&C.

Web Appendix

A More Details of State Cider Regulations

States have separate regulations for beer, wine, and distilled spirits, but hard cider as a new category does not have its own group, so its classification as beer or wine varies across states and the type of regulations. Many states classify alcoholic beverages based on how they are made (e.g., Washington), in which case cider is considered a wine because it is obtained from the fermentation of fruit juice. Some states classify drinks based on alcohol content (e.g., Iowa), and cider with the usual amount of alcohol (5% ABV) is classified as a beer. Even within a state, the exact classification can vary depending on the policy purpose despite its general classification. For example, in Florida, cider producers need to obtain a wine manufacturing license to make cider. But they pay a much lower excise tax (\$0.89 per gallon) than wine producers (\$2.25 per gallon), and the tax is closer to beer excise tax (\$0.48 per gallon). For this paper, I collect the specific regulations cider manufacturers and retailers need to follow for each state.

The 16 Post & Hold states are Connecticut, Delaware, Georgia, Idaho, Indiana, Kansas, Massachusetts, Maine, Michigan, New Hampshire, New Jersey, Oregon, South Dakota, Tennessee, Vermont, and Wyoming. The Post part usually requires filing price schedules at the state office but sometimes requires posting prices online (e.g., Connecticut) or maintaining a price list at the licensed location for inspection (e.g., Oregon). The 14 states that ban volume discounts are Connecticut, Delaware, Idaho, Kansas, Louisiana, Maine, Michigan, North Carolina, Ohio, Oklahoma, Oregon, Tennessee, Vermont, and Wyoming. Washington was a control state before Dec 8, 2011, and satisfied both. There are no changes to the laws discussed in this paper, except for the privatization in Washington state on Dec 8, 2011.

B Robustness Checks

B.1 Relationship of Cider Availability and Manufacturer Past Beer Share by Retailer Size and Retailer Type

Table 10: New Product Availability and Past Beer Share by Retailer Size

	Coefficient (SE)
Mfr. Past Beer Share by Retailer	
Top 20 Chain	0.48*** (0.17)
Below 20 Chain	0.50*** (0.14)
Umbrella Brand Share by Retailer	
Top 20 Chain	6.41*** (1.27)
Below 20 Chain	3.82*** (0.90)
Brand-County-Week FE	Y
Store-Week FE	Y
Observations	10,499,897
R ²	0.70

Notes: Retailer ranking is based on the total dollar sales of cider, 2006-2016. Standard errors are clustered at the brand-chain level. *p<0.1; **p<0.05; ***p<0.01

Table 11: New Product Availability and Past Beer Share by Retailer Type

	Coefficient (SE)
Mfr. Past Beer Share by Retailer	
Convenience Stores	0.31* (0.16)
Drug Stores	0.59*** (0.15)
Grocery Stores	0.44** (0.21)
Liquor Stores	0.71* (0.36)
Mass Merchandisers	0.47** (0.23)
Umbrella Brand Share by Retailer	
Convenience Stores	4.91** (2.36)
Drug Stores	5.14*** (1.46)
Grocery Stores	5.16*** (1.09)
Liquor Stores	-1.53 (1.70)
Mass Merchandisers	8.92*** (1.36)
Brand-County-Week FE	Y
Store-Week FE	Y
Observations	10,499,897
R ²	0.71

Notes: Standard errors are clustered at the brand-chain level. *p<0.1; **p<0.05; ***p<0.01

B.2 Local Logistics Efficiency

Table 12: Brand-Store Pairs with Positive Manufacturer Past Beer Share at the Store

	carry	log(#UPC)
	(1)	(2)
Mfr. Past Beer Share by Retailer	0.38*** (0.15)	1.54*** (0.56)
Umbrella Control?	Yes	Yes
Brand-County-Week FE	Yes	Yes
Store-Week FE	Yes	Yes
Observations	8,028,897	2,903,961
R ²	0.72	0.89

Notes: This table shows that the main results are robust when focusing on brand-store pairs that the store has already been buying beer from the brewer. This rules out the alternative explanation of last-mile shipping efficiency (i.e., stores are just buying cider from manufacturers that they have already been buying beer from). Standard errors are clustered at the brand-chain level. *p<0.1; **p<0.05; ***p<0.01

B.3 Do the Brewers Replace Beer with Cider on the Shelf?

Table 13: Manufacturer Current Beer Share and the Number of Cider Products

	Mfr. Current Beer Share at Store
#UPC, Mfr. Cider at Store	0.00061* (0.00033)
Manufacturer-Store FE	Y
Manufacturer-Week FE	Y
Store-Week FE	Y
Observations	10,728,436
R ²	0.95386

Notes: This table shows that a manufacturer's current beer share at a retailer does not decrease as the retailer carries its cider, suggesting the relationship's effect on distribution is not driven by direct replacement of the manufacturer's beer (product or facing). The result also suggests there is no specific substitution between the same manufacturer's cider and beer. The regression is at the manufacturer-store-week level. The sample includes the periods before and after the cider launches. Standard errors are clustered at the manufacturer-chain level. *p<0.1; **p<0.05; ***p<0.01

B.4 Relationships in States with Strict Beer Wholesale Pricing Rules

Table 14: States with Strict Beer Wholesale Pricing Rules

	carry	log(#UPC)
	(1)	(2)
Mfr. Past Beer Share by Retailer	0.48*** (0.14)	1.01** (0.40)
$\times \mathbf{1}\{\text{Beer P\&H or Ban Q.D.}\}$	-0.004 (0.14)	1.04* (0.61)
Umbrella Control?	Yes	Yes
Brand-County-Week FE	Yes	Yes
Store-Week FE	Yes	Yes
Observations	10,499,897	3,320,277
R ²	0.70	0.88

Notes: This table shows that the main results hold for states with strict beer wholesale pricing rules. This rules out the explanation that brewers offer wholesale discounts for their beer to induce retailers to buy their cider. Standard errors are clustered at the brand-chain level. *p<0.1; **p<0.05; ***p<0.01

B.5 Category Captaincy

Table 15: Regression without Each Retailer’s Highest Share Manufacturer

	carry	log(#UPC)
	(1)	(2)
Mfr. Past Beer Share by Retailer	0.88* (0.47)	3.26* (1.85)
Umbrella Control?	Yes	Yes
Brand-County-Week FE	Yes	Yes
Store-Week FE	Yes	Yes
Observations	5,067,116	2,045,359
R ²	0.83	0.93

Notes: This table shows that the main results are robust without each retailer’s highest share manufacturer. Since nearly half of the observations are removed, the results are not as statistically significant as the main ones. Standard errors are clustered at the brand-chain level. *p<0.1; **p<0.05; ***p<0.01

B.6 Relationships with Chains vs. with Stores

An alternative explanation for the relationship’s effect on distribution is that the relationships are established between stores and manufacturers/wholesalers instead of between retailers and manufacturers. To address this concern, I add the manufacturer’s past beer share at a given store to the regression as an additional control, which proxies the manufacturer-store relationship and captures the variation across stores for a given retailer. If the relationship’s effect is consistent across stores, then it suggests the relationships are established with retailers, and the retailer-level coefficient should capture the effect. By contrast, if the relationships are established with stores, the store-level coefficient should pick up the effect.

Table 16 shows that the relationships are mainly established with retailers instead of stores. The estimates of the retailer-level coefficients are close to the main results, while the estimates of the store-level coefficients are much smaller. Thus, the store-level variable cannot explain away the retailer-level effect and the relationships are established with retailers.

Table 16: Relationships with Chains vs. with Stores

	carry	log(#UPC)
	(1)	(2)
Mfr. Past Beer Share by Retailer	0.49*** (0.14)	1.37*** (0.49)
Mfr. Past Beer Share by Store	-0.06 (0.06)	0.25** (0.11)
Umbrella Control?	Yes	Yes
Brand-County-Week FE	Yes	Yes
Store-Week FE	Yes	Yes
Observations	10,499,897	3,320,277
R ²	0.70	0.88

Notes: Standard errors are clustered at the brand-chain level. *p<0.1; **p<0.05; ***p<0.01

B.7 Distribution Effect in Regulated and Unregulated States

This section examines whether the relationship’s effect on distribution exists only in the unregulated states or all states. In Table 17, I add an interaction term of past beer share and the regulated dummy to the main regression. The results show the distribution effect exists in both types of

Table 17: Relationship Effect and the Type of States the Stores Are in

	carry	log(#UPC)
	(1)	(2)
Mfr. Past Beer Share by Retailer	0.53*** (0.16)	1.44*** (0.51)
× 1{P&H or Ban Q.D.}	-0.23 (0.14)	0.29 (0.56)
Umbrella Control?	Yes	Yes
Brand-County-Week FE	Yes	Yes
Store-Week FE	Yes	Yes
Observations	10,499,897	3,320,277
R ²	0.70	0.88

Notes: Standard errors are clustered at the brand-chain level. *p<0.1; **p<0.05; ***p<0.01

Table 18: Relationship Effect and the Retailer’s Share of Stores in Regulated States

	carry	log(#UPC)
	(1)	(2)
Mfr. Past Beer Share by Retailer	0.44*** (0.15)	1.45*** (0.48)
× Share of Stores in Regulated States	0.08 (0.23)	0.33 (1.01)
Umbrella Control?	Yes	Yes
Brand-County-Week FE	Yes	Yes
Store-Week FE	Yes	Yes
Observations	10,499,897	3,320,277
R ²	0.70	0.88

Notes: Standard errors are clustered at the brand-chain level. *p<0.1; **p<0.05; ***p<0.01

states. The interaction term is statistically insignificant for the binary variable “carry” (-0.23) and the log number of products carried (0.29). Table 18 presents an alternative specification that includes an interaction of past beer share and the retailer’s share of stores in regulated states. The interaction terms are also insignificant. Thus, both results suggest the distribution effect exists in all states.

C Ratios γ_r/γ_m , ω_r/ω_m , γ/ω

The estimates of θ and ρ contain information of the underlying weight parameters ω and KKT multipliers γ . Specifically, the following ratios can be backed out from the model estimates:

$$\frac{\gamma_r}{\gamma_m} = \theta_{mr} \frac{\rho_r}{\rho_m}; \quad \frac{\omega_r}{\omega_m} = \theta_{mr} \frac{1 - \rho_r}{1 - \rho_m}; \quad \frac{\gamma_r}{\omega_r} = \frac{\rho_r}{1 - \rho_r}; \quad \frac{\gamma_m}{\omega_m} = \frac{\rho_m}{1 - \rho_m}.$$

Figure 10 presents the ratios and how they change with the past beer share. If we fix the value of one of $\{\gamma_r, \gamma_m, \omega_r, \omega_m\}$, the other three can be recovered from the ratios.

First, the ratio $\gamma_r/\gamma_m < 1$ (Figure 10 panel (a)) implies that the KKT multipliers are greater for the manufacturers' incentive constraints than the retailers' incentive constraints. In other words, relaxing the manufacturers' incentive constraints will yield a larger increase in the objective (3) than relaxing the retailers'. Although the two estimates shown in the figure are different (based on $\hat{\theta}^w$ and $\hat{\theta}^a$, respectively), they both suggest $\gamma_m > \gamma_r$.

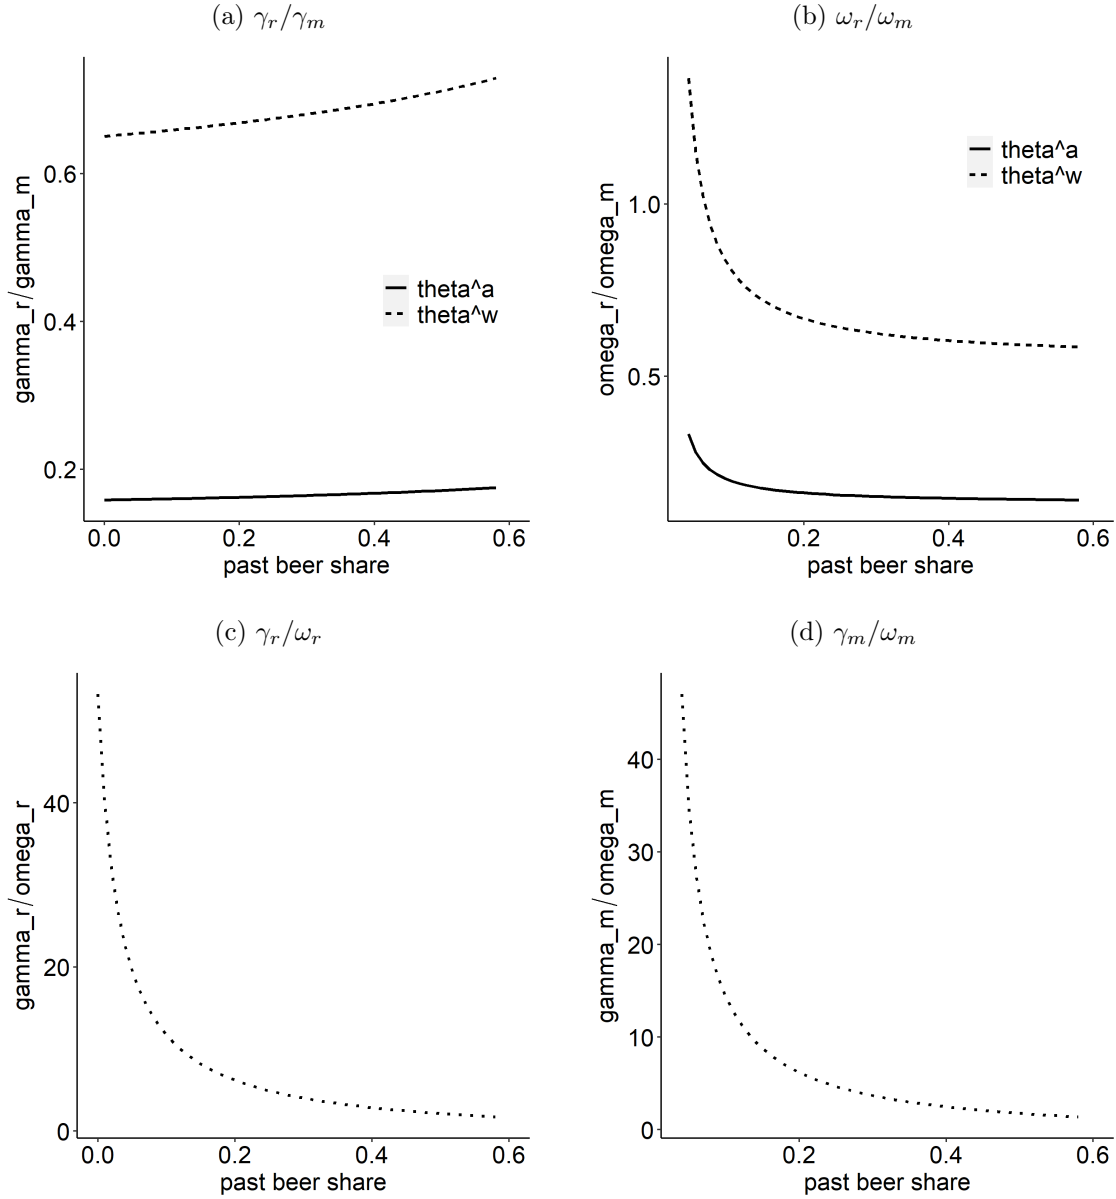
Second, the ratio ω_r/ω_m (Figure 10 panel (b)) shows how manufacturers' and retailers' profits are weighted in the objective (3). We can see the retailers' weights decrease as the past beer share increases. The weights are higher for ABI and MillerCoors than for retailers ($\omega_r/\omega_m < 1$). Note that a retailer's profits are usually higher than a manufacturer's profits per store. Thus, percentage-wise, the retailers' profits could be larger weighted than the manufacturers' profits (e.g., a 10% increase in r 's profits might be preferred over a 10% increase in m 's profits).

Third, the ratios γ_r/ω_r and γ_m/ω_m (Figure 10 panels (c) and (d)) capture the gain from relaxing the incentive constraints normalized by the level of profit weights. We can see as the past beer share increases, the normalized gain decreases. Because the gain from relaxing constraints should be negatively related to the coordination level, the pattern is consistent with the finding that trust and coordination increase in the past beer share.

D Profit-Weight Approach

In this section, I present the results of the profit-weight approach. Compared to the repeated game-based model in the main paper, the profit-weight approach assumes $\rho = 0$ (equations (10), (11), (19), and (22)). In other words, a manufacturer or a retailer will maximize a weighted sum of its

Figure 10: Ratios of γ and ω



own profits and its partner's profits. Despite not being generally equivalent to a repeated game, the profit-weight model has a simple and clear interpretation that the weight parameter captures the extent to which the two firms incorporate each other's profits in their decision. The profit-weight approach generates very similar counterfactual results as the full repeated game-based approach.

Table 19 and Table 20 present the estimates for the wholesale pricing stage and the assortment choice stage, respectively. They demonstrate that the extent to which two firms incorporate each

Table 19: Pricing Estimates, Profit-Weight Approach

θ_1^w	0.41*** (0.07)
Average retail margin	35.9%
Average wholesale margin	35.3%
Average wholesale price (6-pack)	\$6.26

Notes: Standard errors are clustered at the product-store level. *p<0.1; **p<0.05; ***p<0.01

Table 20: Assortment Choices Estimates, Profit-Weight Approach

θ_0^a	-0.22*** (0.008)
θ_1^a	2.20*** (0.07)
σ^ν	0.32*** (0.003)

Notes: Standard errors in parentheses are based on the assortment stage only. *p<0.1; **p<0.05; ***p<0.01

other's profits increases in the past beer share. Moreover, Table 21 reports the pair-specific estimates for the assortment stage. Although these estimates could contain a lot of noise, they nevertheless show an interesting heterogeneity of relationships across retailers. As we can see, some retailers have a very close relationship with a certain brewer (e.g., retailer 79), while some do not have good relationships with any brewers (e.g., retailer 236).

Table 21: Pair-Specific Assortment Choices Estimates, Profit-Weight Approach

Retailer	Relationship, θ_{mr}^a				$\sigma_r^\omega/100$
	ABI	MillerCoors	Heineken	BostonBeer	
130	0.89 (0.05)	0.52 (0.06)	0.12 (0.04)	-0.14 (0.02)	0.22 (0)
32	0.21 (0.04)	-0.6 (0.03)	-0.26 (0.03)	-0.27 (0.02)	0.23 (0)
158	0.61 (0.14)	0.68 (0.09)	-0.27 (0.07)	-0.45 (0.03)	0.35 (0.01)
79	1.66 (0.08)	0.08 (0.1)	-0.49 (0.07)	-0.27 (0.02)	0.18 (0.01)
6901	1.49 (0.08)	1.86 (0.12)	-0.45 (0.06)	-0.18 (0.02)	0.22 (0.01)
9	3.62 (0.45)	3.44 (0.48)	1.87 (0.41)	1.07 (0.08)	1.42 (0.06)
111	0.21 (0.07)	0.29 (0.07)	0.06 (0.05)	-0.33 (0.03)	0.19 (0.01)
117	0.72 (0.12)	0.12 (0.1)	0.32 (0.09)	-0.5 (0.04)	0.14 (0.01)
843	0.52 (0.11)	0.1 (0.12)	-0.58 (0.1)	-0.57 (0.02)	0.19 (0.01)
236	-0.49 (0.02)	-0.64 (0.03)	-0.71 (0.02)	-0.62 (0.01)	0.07 (0)
184	0.33 (0.05)	0.16 (0.09)	-0.04 (0.05)	-0.42 (0.02)	0.14 (0)
89	-2.11 (0.51)	0.41 (1.18)	1.32 (0.29)	-0.47 (0.05)	0.3 (0.03)
182	0.6 (0.12)	0.71 (0.09)	0.68 (0.1)	-0.47 (0.02)	0.27 (0.01)
248	0.86 (0.1)	0.36 (0.13)	0.46 (0.08)	-0.49 (0.03)	0.13 (0.01)
4904	-0.24 (0.22)	0.1 (0.28)	0.01 (0.35)	0.49 (0.06)	0.14 (0.01)
6	0.13 (0.08)	0.64 (0.12)	0.44 (0.11)	-0.57 (0.03)	0.14 (0.01)
128	0.19 (0.09)	1.77 (0.18)	-0.35 (0.06)	-0.31 (0.05)	0.17 (0.01)
90	0.57 (0.08)	-1.7 (0.12)	-1.03 (0.04)	-0.61 (0.02)	0.1 (0)
34	-0.36 (0.12)	-1.19 (0.13)	-2.77 (0.16)	-0.45 (0.05)	0.14 (0.01)
210	3.43 (0.23)	3.99 (0.29)	3.42 (0.22)	0.98 (0.1)	0.41 (0.02)
59	0.84 (0.35)	0.5 (0.3)	-1.38 (0.31)	-0.66 (0.04)	0.61 (0.07)

844	0.98 (0.13)	0.53 (0.18)	0.05 (0.1)	-0.32 (0.04)	0.12 (0.01)
194	0.26 (0.16)	1.7 (0.2)	0.56 (0.12)	-0.27 (0.05)	0.11 (0.01)
199	0.6 (0.14)	0.24 (0.17)	0.67 (0.17)	-0.01 (0.04)	0.21 (0.01)
69	-0.18 (0.37)	1.29 (0.78)	-0.89 (0.28)	-0.36 (0.07)	0.34 (0.03)
212	-0.55 (0.08)	-7.76 (1.21)	-1.47 (0.06)	-0.79 (0.02)	0.09 (0)
257	-0.93 (0.28)	-0.95 (0.15)	-1.16 (0.23)	-0.85 (0.05)	0.2 (0.03)
97	0.07 (0.15)	-0.14 (0.17)	-0.72 (0.16)	-0.43 (0.04)	0.1 (0.01)
50	0.12 (0.06)	0.19 (0.12)	0.69 (0.11)	-0.29 (0.03)	0.07 (0)
174	0.95 (0.39)	0.11 (0.33)	-0.32 (0.22)	-0.3 (0.09)	0.17 (0.02)
36	0.02 (0.26)	0.64 (0.44)	0.35 (0.29)	-0.09 (0.1)	0.24 (0.03)
62	0.63 (0.2)	0.05 (0.18)	-0.52 (0.12)	-0.62 (0.05)	0.11 (0.01)
869	-0.06 (0.14)	0.57 (0.17)	-0.13 (0.1)	-0.11 (0.1)	0.11 (0.01)
185	-0.04 (0.08)	0.13 (0.12)	-0.55 (0.04)	-0.14 (0.06)	0.07 (0)
136	2.78 (0.47)	4.36 (1.85)	0.54 (0.33)	-0.28 (0.11)	0.14 (0.02)
315	2.87 (0.46)	3.87 (0.59)	0.58 (0.41)	1.54 (0.3)	0.17 (0.02)
61	14.6 (0.45)	4.72 (0.19)	-0.15 (0.31)	-0.35 (0.03)	0.05 (0.01)
221	0.8 (0.97)	0.79 (0.79)	0.02 (0.46)	-0.45 (0.11)	0.8 (0.16)
123	-1.03 (0.18)	-0.81 (0.14)	0.92 (0.22)	-0.32 (0.07)	0.13 (0.01)
889	-0.19 (0.3)	-0.01 (0.32)	-7.1 (2.02)	-0.34 (0.09)	0.23 (0.03)
328	-1.02 (0.2)	0.17 (0.11)	-0.86 (0.05)	-0.53 (0.03)	0.04 (0)
2	0.4 (0.18)	0.49 (0.32)	-0.68 (0.16)	-0.56 (0.07)	0.06 (0.01)
770	-0.5 (0.6)	0.87 (0.86)	-4.55 (1.03)	0.23 (0.29)	0.27 (0.04)

Notes: Retailers ordered by total cider sales. Standard errors based on the assortment stage only.