Title

Essays in Industrial Organization

Author

Perez, Daniel Reyes

Publication Date

2024

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

Essays in Industrial Organization

A dissertation submitted in partial satisfaction of the requirements for the degree Doctor of Philosophy in Economics

by

Daniel Reyes Perez

2024

ABSTRACT OF THE DISSERTATION

Essays in Industrial Organization

by

Daniel Reyes Perez

Doctor of Philosophy in Economics
University of California, Los Angeles, 2024
Professor John William Asker, Chair

This dissertation examines questions in industrial organization regarding the evaluation of antitrust and procurement policies. In the first two chapters, I examine whether antitrust agencies should consider the effects of product variety changes due to a merger. In the first chapter I test whether the merger had any effect on product variety directly. I find evidence of product variety changes due to a merger in the context of the MillerCoors merger of 2008. I find that the merged firm decreased the number of brands offered and offset this by increasing product variety in more successful brands. However, under a difference-and-differences framework, I find that product variety declined relative to that of other top competitors.

In the second chapter, I examine the welfare effects of product variety changes and compare them to the welfare effects of prices to see what should be the priority of antitrust agencies. I estimate demand for the MillerCoors merger in the postmerger period, expanding on previous work in the literature. In a set of two counterfactuals, I test the value of new products created after the merger and the value of discontinued products lost after the merger. I find that the merger increased consumer surplus from changes in product variety. Benchmarking this to results in the literature, I find that the effects of product addition

and discontinuation are approximately 34% and -4%, respectively, of the consumer welfare effects of the postmerger price changes in the presence of coordinated pricing found in prior work.

In the third chapter, I provide a theoretical framework on how restrictions on firms participating in government procurement auctions affect local government procurement costs, firm participation for procurement decisions and firm investment decisions under the context of the Buy America Policy change in 2015. This policy increased the domestic materials requirement for federal funds to be used for purchases of transit goods, thereby restricting foreign firms from participating in procurement and potentially raising the costs of domestic firms that rely on foreign capital for their products. I find federal transportation funds used for railcars increased while funds used for buses remained flat. I then propose three models to examine the theoretical impact of the policy. The first model estimates a standard first price sealed-bid auction model standard for procurement. The second model separates bidders into two types: one domestic type which draws from a higher mean cost distribution, and one foreign type which draws from a lower mean cost distribution. I find costs and markups increase, while the participation of foreign firms declines.

The dissertation of Daniel Reyes Perez is approved.

Michela Giorcelli

William M. Rafey

Hugo Andres Hopenhayn

John William Asker, Committee Chair

University of California, Los Angeles

2024

To my mother, for caring and supporting me every day of the past six years.

To my brother, the funniest and smartest one in the family.

To my sister, for being a constant source of joy through her creativity.

To my father, for being the biggest source of inspiration.

Para mi abuela, por su amor y apoyo. Te quiero mucho.

Contents

List of Figures

List of Tables

ACKNOWLEDGMENTS

I would like to acknowledge the head of my advising committee, John Asker, and the rest of the committee members: Will Rafey, Michela Giorcelli, and Hugo Hopenhayn. Thank you to Will for being a friendly mentor on my projects, and consistently prodding me to push them further to contribute more to the field. Thank you to Michela for her advice on my projects within the dissertation and on projects beyond my field. Thank you to Hugo for assisting me on the theoretical models of the first and second chapter, which later helped me become a better research communicator. Finally, thank you to John for your consistent advice and patience throughout this dissertation, as it would not be complete without your thorough feedback and help. I appreciate you always willing to step in with support whenever I faced a roadblock, and your lessons will continue to inspire me as an economist.

I would also like to acknowledge the colleagues and professors that helped improve the draft. Thank you to Martin Hackmann for providing invaluable advice on the first chapter and for providing a critical eye to improve many sections of the dissertation. Thank you to the rest of the Industrial Organization faculty at UCLA for their assistance in seminars, in one-on-one meetings and for further honing the dissertation and my research skills. Thank you to my colleagues in the Industrial Organization field at UCLA: Domenico Fabrazi, Chandni Raja and Andrew Heinzman. Our many IO lunches helped me form my research paper and created a consistent community that I hope continues well past our times at UCLA. I am honored to be at the same institution as you all and to grow as researchers together. Thank you also to the other colleagues in my UCLA cohort, such as Manu Navjeevan, Calvin Kuo, Danny Ober-Reynolds, Nicole Gorton Caratelli and Jesper Böjeryd, for providing friendship, camaraderie and research advice throughout the years.

I would also like to acknowledge the staff and financial support received throughout the program. Thank you to the Economics Department office for their help throughout the Ph.D. on logistics and administrative tasks, and being prompt about any minute question I may have. Thank you to financial support from UCLA Graduate Division, the Economics Department, and from John Asker and Michela Giorcelli for this dissertation. Finally, I would like to thank staff at IRI and BMC for assistance with datasets used in the first and second chapter of the dissertation.

Finally, thank you to my friends and family who helped me through the six years of the Ph.D. On days when I had no will to continue on, your support and hopes kept me going until I could move forward with confidence. I would especially like to thank my parents, my sister, my brother, my abuelita, my cat, and my wonderful extended family, which I will refrain from writing every name of in order to keep the acknowledgment section shorter than the dissertation itself. From large gestures to small check-ins, for those still in our lives and those who have passed on, this dissertation would not have been possible without your support. I dedicate this work to you all.

VITA

Daniel Reyes Perez

Education

University of California, Los Angeles

M.A. in Economics 2018-2020

Occidental College

B.A. in Economics and Mathematics 2012-2016

Research in Progress

"How Much Do Consumers Care About New and Discontinued Products? A Case Study of the MillerCoors Merger" (*Job Market Paper*)

"The Effects of Domestic Restrictions in Government Procurement for Transportation Goods"

Relevant Positions

University of California Los Angeles (UCLA)

Research Assistant 2019-2021

Federal Reserve Bank of Kansas City

Research Associate 2016-2018

Fellowships, Honors, and Awards

Fellowship, Department of Economics, UCLA. 2024

Graduate Summer Research Mentorship, UCLA. 2022

Cota Robles Fellowship, UCLA. 2018 - 2023

Research funds and grants

Lewis L. Clarke Graduate Fellowship Fund, UCLA. 2022

Presentations

International Industrial Organization Conference 2024

Chapter 1

Measuring the impact of the MillerCoors Merger on Product Varieties

1.1 Introduction

One issue that regulators face in merger analysis is how to quantify the nonprice effects of a merger. According to a recent OECD report summarizing the policies of twenty-one competition agencies around the world, eighteen explicitly have policies addressing nonprice effects of mergers, but the majority do not address these issues unless there are "claims made by merging parties, their customers/consumers, and rivals".[1] However, it is unclear how one such nonprice effect, that on product variety, should be weighted in the social optimum, as some agencies consider them second order to price effects while others consider them of equal importance to price effects. Secondly, it is unclear whether they occur in mergers. One aspect of resolving these issues requires quantifying the effect of the merger on product

[1]In such cases, the complaints may be addressed through qualitative evidence in the absence of direct measurement (Capobianco 2018).

variety.[2]

In this chapter, I use the setting of the brewery industry after the MillerCoors merger of 2008 to document changes in product variety due to a merger. Prior work has not considered product variety changes and their potential effects on consumers in the context of this widely studied merger. I first describe the setting, show how price is limited in this market and discuss how much the product variety effects can be attributed to the merger itself. I find that product variety fell relative to that of competitors.

I focus on the brewing industry because it is one such industry where heterogeneous products are offered and product variety is valuable for consumers.[3] As described in detail by Miller and Weinberg (2017) and Weinberg, Sheu, and Miller (2019), brewing firms often priced their products very similarly to their competitors', especially for brands owned by Miller, Coors and Anheuser-Busch Inbev. Several reasons for this similarity were given, such as implicit price collusion, the role of retailers and distributors in price setting and the role of competitors in hampering such implicit collusion through price undercutting. Because of this, firms may decide to cut costs through removing products or capture more profit by introducing new products outside of their flagship brands.

The main data source for this chapter is the IRI Marketing dataset, which provides data on the demand for beer between 2002 to 2012. Using this dataset, I am able to observe in each market which products were removed and which were added before and after the MillerCoors merger of 2008. This merger combined the second and third largest brewery companies in the United States, leading to exogenous changes in costs and market structure. Thus, product varieties changed at the market level, with both product additions and discontinuations after the merger. I additionally supplement these data with the Beverage Marketing Company (BMC) database, which has producer and distributor characteristics,

[2]The second chapter of the dissertation examines the other aspect, measuring the consumer welfare effect of product variety changes for direct comparison.

[3]More detail on why firms may compete on product variety, rather than prices in this market, can be found in Chapter 2.

to provide added context on the market.

I first examine whether product variety, measured in terms of both individual products and brands more generally defined, changed after the merger and whether this can be attributed to the merger itself. Under a basic linear model, I find that, while the merger had a minor negative effect on the number of brands provided by both Miller and Coors, it had no effect on the number of products supplied by the merged company. This provides evidence of the phenomenon described in Atalay et al. (2023) of the merged firm focusing on core products rather than new products or products at the periphery of its business.

However, brand variety in the merged firm declined relative to that of its competitors. I use a difference-in-differences design, comparing brand variety between Miller and Coors and the top firms in terms of the market share measure commonly used in the literature, and find that brand variety at the merged company declined relative to its competitors' by 22%. This decline holds when I detrend the data of premerger time trends, improving elements of the original difference-in-differences design. This negative effect relative to competitors is robust to other robustness checks conducted in the appendix.

1.1.1 Literature Review

This chapter contributes to two main strands of literature: the stream on the effects of changing product variety and that on the MillerCoors merger of 2008. I contribute to each by providing a framework for how to model changing product variety and estimating this effect for the merger.

This chapter contributes to the growing literature on product variety changes. As mentioned previously, Wollmann (2018) studies the impact of product variety changes on consumers under an exogenous shock to product variety through bailouts in 2007. Similarly, Fan and Yang (2022) study how firms may reallocate products after an acquisition of a craft brewery. I make one important contribution complimenting their work in this chapter: my use of different data. Due to data limitations, the authors of the aforementioned works cannot identify which breweries or beers drive their results. My data are not subject to this limitation, although they cover fewer markets and a shorter time frame. Nevertheless, using these less restricted data allows me to provide more information on firm specifics. By avoiding assumptions on what products leave the market and enter, I provide an empirical test of the conclusions from these authors' models.

In addition, I contribute to the literature on the MillerCoors merger of 2008 and the implications of the changes in the beer market occurring around this time. Ashenfelter, Hosken, and Weinberg (2015) examine the impacts of the merger on prices and transportation costs for the firm and their relative importance. They ultimately find that the price increases were offset by efficiency gains from the merger, leading to very little change in prices attributable to the merger itself. In other work, Miller and Weinberg (2017) examine the price effects of the merger and find that they are much higher than predicted by a model capturing the merger's effect in facilitating price coordination among top firms. In a follow-up paper, Weinberg, Sheu, and Miller (2019) find evidence of tacit pricing coordination after the merger in the brewing industry that potentially explains these effects. Related to merger analysis, Khmelnitskaya, Marshall and Orr (2024) measure the impacts of scale and scope economies for the MillerCoors merger and how these can affect pricing decisions post-merger. I examine product variety effects in this merger, a first for studies regarding the MillerCoors merger of 2008.

My results both validate and provide additional context to results previously seen in the literature. Examining the merger's effect on product variety, I find results similar to those of Atalay et al. (2023), who find, using a large-scale event study, a slight decline in product variety from a merger. They document the phenomenon of brand consolidation: firms cutting back their product variety in certain brand lines to focus on their highest-revenue products in other brand lines. I observe this phenomenon in my study as well. After the merger, the number of brands falls within each market relative to other top competitors.

The paper proceeds as follows. In section 2, I describe the setting of the brewing industry, the merger itself, and the suitability of this merger and setting for testing the welfare effects of product variety and price changes. In section 3, I describe the data and provide key summary statistics for each dataset. In section 4, I describe the reduced-form models and results to examine the merger's impact on product variety. Finally, in section 5, I conclude.

1.2 Setting

In this section, I expand on the features of the brewing industry that make it an ideal setting to examine this relationship's effect on consumer welfare. I additionally discuss the MillerCoors merger of 2008, the main exogenous shock to market structure that is the focal point of this analysis.

1.2.1 The U.S. Commercial Beer Industry

My research setting is the United States commercial beer industry, a market dominated by fifteen firms with large market share, which allows strategic interactions among them in product variety choices. I choose this industry because it is characterized by easily identifiable product differentiation via brands, packaging type and size. Additionally, there is evidence of limited price changes within this market. This industry was largely stable and faced few aggregate shocks to product variety until 2008, when the Miller-Coors merger occurred; this quasi-experimental setting lets me focus on product variety.

The United States beer industry shares many similarities with other branded consumer product industries; however, there are some important differences that I use to my advantage in this study. Similarly to firms in other branded consumer product industries, beer companies compete on prices and further improve their competitive stance in the industry through quality, product introduction, advertising and sales. The major difference is that distributors are a vital and highly regulated section of the market. By law, brewers first sell to distributors, and distributors then sell to stores. Distributors market various types of beverages, such as beers, alcohols, soft drinks and others.[4] Several states have enacted additional laws and, in some cases, restrictions on distribution. For example, each state imposes its own excise tax on beer distributors in the state, ranging from \$ 0.02 per gallon in Wyoming to \$ 1.29 per gallon in Tennessee. Some states, such as New York and New Jersey, also impose an excise tax when manufacturers ship products to the state.[5] These laws provide variation in prices and distribution, which influences beer consumption at the local level.

The presence of the distribution market makes distance and negotiation costs much more pertinent for product variety in this than in other industries. Because brewers must first sell to distributors, who then sell to consumer-facing businesses, distribution costs increase final costs and therefore final product variety. Distribution costs can take two forms that I focus on in this chapter: distance costs and negotiation costs. For the former, costs rise with brewers' distance from distributors, cutting into individual product profits. For the latter, brewery companies must negotiate what they pay distributors for distributing brewery products.[6] As will be described later, these two types of costs were the main impetus for the merger between Miller and Coors.

Product variety is also tied closely to individual consumer preferences, which are fairly strong in the U.S. commercial beer industry. Typically, consumers have strong preferences for beer produced within their region. For example, Anheuser-Busch is the market leader in

[4] See the National Beer Wholesalers Association (NBWA) website for more information: https://www.nbwa.org/about/what-beer-distributor.

[5] A discussion on 2021 excise tax trends can be found here: https://taxfoundation.org/excise-taxes-excise-tax-trends/#Alcohol.

[6] This can lead to exclusive deals between breweries and distributors, such as the exclusivity bonuses provided by Anheuser-Busch Inc. More information is accessible at:https://www.bizjournals.com/stlouis/stories/2008/03/31/daily73.html.

St. Louis, where its central brewery is located. This market feature allows plausible consumer variation in preferences as well. While firms still have other tools to increase their market share, such as advertising and temporary sales,[7] trends in these components were largely stable before the merger. Overall mean market shares remained stable around 18% prior to the merger, and market shares in each region also held steady prior to the merger.

Given the role of distributors and these strong consumer preferences, product variety is a vital part of breweries' competitive market. Since 1979, there has been a stark increase in the number of craft beers, with the trend accelerating in the mid-2000s (Elzinga, Tremblay, and Tremblay 2015). Several reasons for this change have been proposed, including deregulation in the industry[8] and local and state policies increasingly friendly to craft brewing (Barajas, Boeing, and Wartell 2017). This increase has led to greater local and national variety for consumers on various product dimensions. In this context, the craft brewing industry has steadily grown, totaling "$26.8 billion, and now account[ing] for just under 27% of the $100 billion U.S. beer market" (Association 2022). This changing craft brewery panorama and the increasing product variety provide the setting for the Coors and Miller merger of 2008.

Finally, prior to the merger, this market was highly concentrated among fifteen firms. While the IRI dataset contains 852 breweries, the 15 largest breweries fell within the top 5 percentile of total market share.[9] On average, the top 15 firms had a combined market share of 94% prior to the merger and 92.9% after the merger. Miller and Coors accounted for nearly a third of this market share, with the two combined having a market share of 29.2% prior to and 27.2% after the merger. Their main competitor and the market leader, Anheuser-Busch InBev, had a market share of 39.9% prior to and 35.6% after the merger.

[7] See Chandra and Weinberg (2018) for a discussion on the role of advertising in this market and within the MillerCoors merger.

[8] See "International Beer Day," accessible at: https://balloon-juice.com/2010/08/05/international-beer-day/

[9] These include Anheuser-Busch, InBev, Anheuser-Busch InBev, SabMiller, Molson Coors, Heineken USA Inc., Groupo Modelo, Boston Beer Co., Cerveceria Costa Rica SA, KPS Partners, Great Lakes Brewing Co., Labatt USA, S&P Company, Constellation Brands Inc., and D.G. Yuengling & Sons Inc.

This concentrated market, as well as the presence of a strong market leader, contributed to Miller and Coors's decision to merge and the rationale for the merger approval by regulatory agencies.

1.2.2 The MillerCoors Merger of 2008

The MillerCoors merger of 2008 combined the United States' second and third largest breweries, respectively. The merger's main motivation was to integrate the geographically distinct production facilities to reduce costs and prices for consumers (see Martin (2007)). I leverage the associated distribution costs changes and the lack of aggregate shocks to argue that the merger could have affected product variety as well.

The MillerCoors merger was announced in October 2007 and finalized on July 1st, 2008.[10] It generated an estimated $ 500 million in cost savings from improving economies of scale (Martin 2007).[11] It was uncertain ex ante whether the merger would pass regulatory scrutiny given market power concerns. First, the merger context was a fairly concentrated market, with Herfindahl–Hirschman index (HHI) estimates for the commercial beer market ranging from 2000 (Ashenfelter, Hosken, and Weinberg (2015)) to 4000 (Tremblay et al. (2005)). Using the IRI dataset, this chapter estimates a national HHI over the entire sample of approximately 2100. Second, past mergers in the beer industry often fell under heavy scrutiny. Prior to this merger, sixteen mergers within the beer industry either had been denied or had seen the merging firms subjected to significant behavioral requirements. Third, many of the products were close substitutes, as Miller and Coors both competed in the four main beer categories. Third, some consulting firms predicted that there would be issues with the merger, with one representative from beverage consultancy Bevmark arguing that the

[10]These two major alcoholic beverage companies existed concurrently until 2020, when the division was later restructured and named the Molson Coors Beverage Company.

[11]In its merger announcement, the company declared the merger would result in an estimated $50 million in savings for the first year, $ 350 million in the second year, and $100 million in the third year. See http://media.corporate-ir.net/media_files/irol/10/101929/molson1.pdf/

merger would lead to "less selection and probably higher prices" (Martin 2007).

Ultimately, however, regulators decided not to challenge the merger. In the review of its investigation, the Department of Justice came to several conclusions. First, it found that the cost savings stated by the company were substantial enough to benefit consumers. Through the integration of geographically distinct production facilities, distribution costs would "be reduced considerably" (Heyer et al. 2009). Transportation cost savings had been the focus of a consulting report prior to the merger, and the Justice Department found no issues with this analysis. Second, the Justice Department found that Miller and Coors competed with each other less than with Anheuser-Busch, which held the highest market share by a significant margin. Finally, it found that the merger was unlikely to increase coordination between firms. For these reasons, the merger was approved and was formally completed on July 1, 2008.

However, later research has found that price coordination may have occurred. Miller and Weinberg (2017) find that the price increases postmerger cannot be fully explained by a transition of Nash equilibriums, and they relate this to evidence that there was price coordination. In follow-up work, Weinberg, Sheu, and Miller (2019) examine the role of implicit price collusion postmerger. They use a price leadership model, where the largest firm announces its price and the rest of the oligopoly sets the price based on this announcement. Here, they find price increases more in line with the actual price increases after the merger. Notably for this study, this price leadership model continues to characterize the postmerger environment, and price leadership is in fact easier to implement with the disappearance of one competitor.

I focus on product variety within the context of this merger for two reasons. First, the merger led to large variation in distribution costs across markets. This was, indeed, the main reason for the merger's occurrence. In its announcement concluding its antitrust investigation, the Department of Justice stated, "In one of the key parts of the investigation, the Division verified that the joint venture is likely to produce substantial and credible savings that will significantly reduce the companies' costs of producing and distributing beer" (Division 2008). These savings largely came from the expansion of Coors' and Miller's production facilities within the United States. Prior to the merger, Coors had two production facilities open: one in Golden, Colorado, and one in Elkton, Virginia. Miller had six, all near large markets such as Irwindale, California (near the Los Angeles metropolitain area), and Fort Worth, Texas.[12] The combined eight breweries would have significantly lower costs of shipping Coors products, which regulators believed could be beneficial to consumers.

The second reason for my product variety focus is that the two firms were constrained in changing prices for the aforementioned reasons. Due to the coordination on prices as documented by Miller and Weinberg (2017), MillerCoors could not lower prices more than its main competitor, Anheuser-Busch InBev. Second, because of the role of retailers, which may price beers similarly for final purchase, any decision to change prices significantly could backfire at this final step. Therefore, firms are restricted in what prices they can offer. With the addition of the distributor market, which further limits price changes upstream, firms' main outlets to affect market share come in the form of new product releases and advertising.[13]

I find evidence of this collusion in prices between 2006 to 2011 in my data, as well. Figure 1.1 shows the differences in mean prices between MillerCoors and Anheuser-Busch InBev products and between MillerCoors products and those of all other top-selling firms. In this period, there was a large divergence in mean prices, peaking at nearly 60 cents in 2008. Despite using a different sample from Miller and Weinberg's (2017), I find that their trend still holds in this expanded dataset in terms of the per-ounce price.

[12] A full list of the breweries can be found on the MolsonCoors blog: `https://www.molsoncoorsblog.com/features/quick-look-my-8-breweries-yes-8`.

[13] While advertising is not the focus of this chapter, a discussion of advertising in this merger appears in Chandra and Weinberg (2018).

Figure 1.1: Differences in Mean Prices Over Time

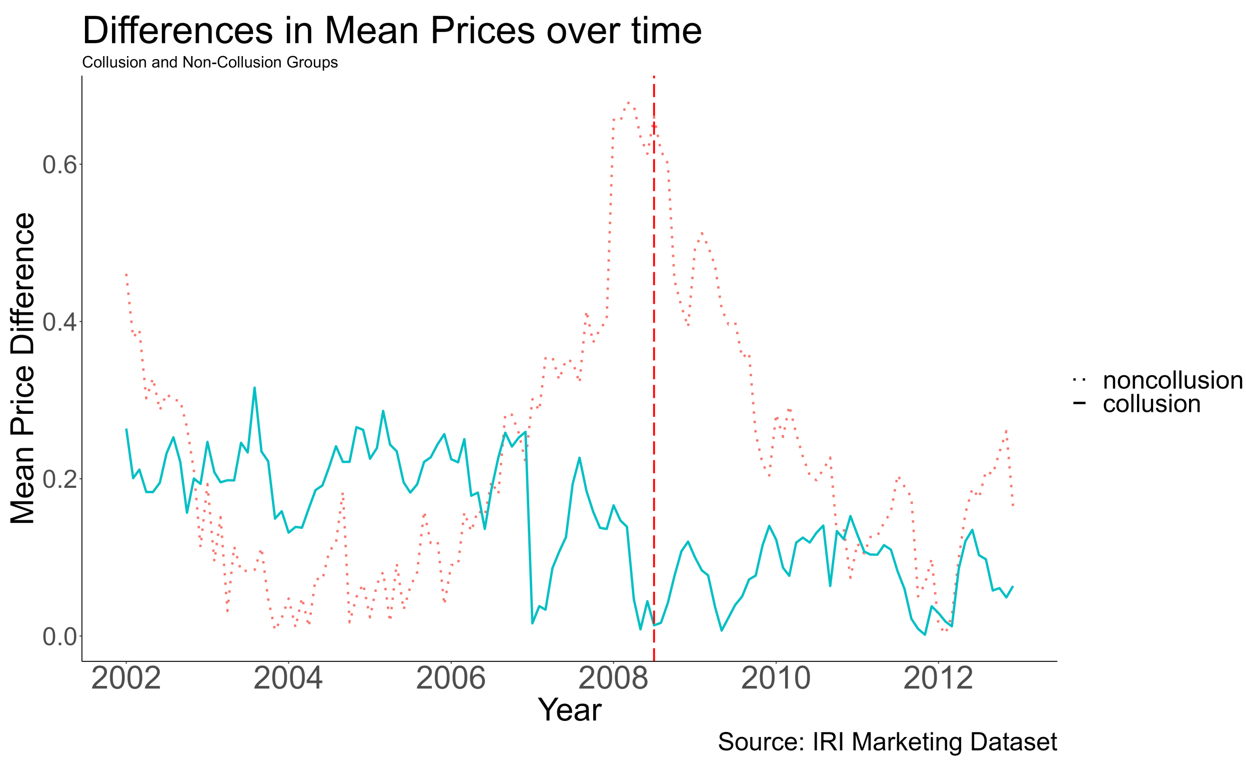

Note: The figure above shows the difference in mean prices between brands from MillerCoors and Anheuser-Busch Inbev, which engaged in implicit collusion on prices according to Miller and Weinberg (2017), and the difference in mean prices between MillerCoors brands and those of all other firms. The sample is limited to the top 5 percentile of firms in national market share. Data are from the IRI dataset.

1.3 Data

Our main datasets, the IRI and BMC datasets, cover the consumer and producer sides of the market, respectively. I now present some summary statistics to provide a baseline for future sections. These two datasets allow me to examine the entire market, rather than a subsection, and give me information on production factors to further improve the model estimates. Although the majority of the analysis uses the IRI dataset, I provide more information on the BMC dataset in the appendix. More specific information on the datasets and features of products discontinued or newly added after the merger can be found in the data appendix.

1.3.1 Consumer Level: IRI Dataset

The IRI dataset provides information on consumer-level demand through scanner data, which show what products consumers buy in stores. I use monthly data from thirty-nine metropolitan statistical areas (MSAs) from 2002 to 2012, which allow me to measure product variety at the final good level and observe revenues, market shares and prices.

The beer industry is a branded consumer product industry, and therefore, varieties, prices and quantities can be measured through supermarket transaction data. The IRI marketing dataset spans 2001–2012 and contains anonymized supermarket transaction data from 51 marketing regions. These marketing regions are typically groups of counties, with some regions crossing state lines. Each observation is an individual sale of a product with a unique UPC identifying a product based on the brand, packaging medium, and size in ounces.

I make several changes to the raw data to facilitate estimation and remove markets with unique legislation restricting the representativeness of observations or impeding my ability to define a market. First, there were several major store-level mergers that occurred in 2001 that affect some of the store-level controls that will prove important for estimation purposes.

Therefore, data from the year 2001 are dropped. To better match against state-level data, I do not include data on markets in states that place restrictions on product variety or distribution. Such laws include those limiting the alcohol content of beer sold or prohibiting sales of beer in supermarkets. This criterion leads me to remove eight markets from the dataset. Finally, for estimation purposes and because of unclear definitions of markets, I do not include four markets that consist of entire states. This provides a total of 39 markets in a total of 28 states. For the main demand specification, I additionally subset the brands into the top ten brands by market share. Prior works, such as Miller and Weinberg (2017), Weinberg, Sheu, and Miller (2019), and Ashenfelter, Hosken, and Weinberg (2014), focus on these firms as well. This is partially for computational reasons, but these companies can better be thought of as Miller's and Coors's closest competitors than can local or craft beer brands.

Overall, market shares vary greatly between markets. Figure 1.2 shows the mean market shares of Miller and Coors prior to the merger and of MillerCoors right after the merger and at the end of the sample period. The shares range from less than 3% to nearly 30% for Coors and from less than 1% to over 50% for Miller. The merged company reaches market shares similar to Miller's, ranging from near 0 to over 50% market share.

1.3.2 Producer Level: BMC Dataset

The BMC dataset provides information on brewery and distributor supplier characteristics between 2006 and 2010. I use annual data from this source to provide more information on potential supply-side characteristics that could affect product variety during the merger. The dataset for distributors includes address, type of importer, parent company of products distributed, number of employees, number of trucks, total sales in that year, and region served. The dataset for brewers includes address, capacity, number of employees, number of lines of canned beverages, number of lines of bottled cold and hot beverages, and region

Figure 1.2: Market Shares for Miller and Coors

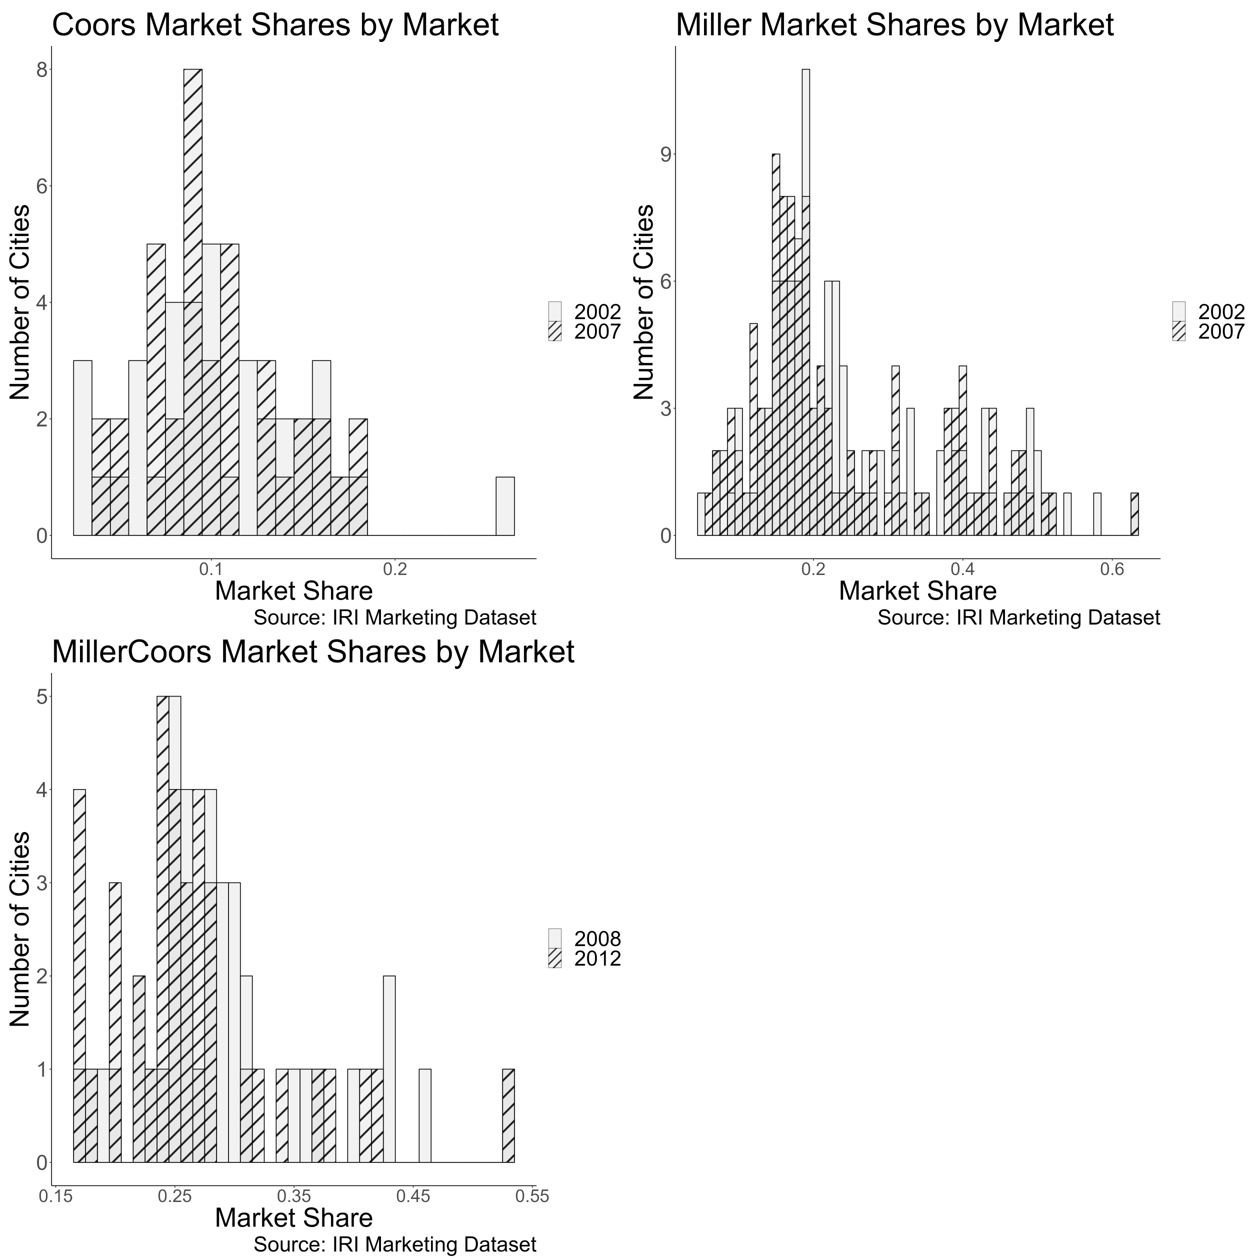

Note: These histograms depict the market share for each market for Coors and Miller for 2002 and 2007 and the market share for MillerCoors for 2008 and 2012. Each observation is a market's annual share for the respective company. Data are compiled through the IRI marketing dataset.

served.

These data sources have several key features that improve the quality of this study. First, the IRI dataset ends four years after the merger, allowing good coverage of the postmerger outcome trends that helps me verify the impacts of the merger. Combining both the BMI and IRI datasets does shrink the postmerger study period to two years, although this still allows me to study the short-term effects of the merger. Second, the ounce-size and packaging medium data allow me to categorize products at a granular level. Due to this, products are classified at the brand–size–packaging level. Finally, the data contain information on the parent company, vendor, and brand linked together. Therefore, although there are many beer brands, all brands can be identified and tracked through time regardless of parent company or vendor changes. I use this dataset to supplement the existing IRI dataset and improve the quality of tracking of brands over time.

A map of the markets as of 2022 and where the breweries are located is show in Figure 1.3. There is some geographic dispersion of both the markets and the breweries, especially for Coors breweries. There are only two Coors breweries: one in Golden, Colorado, and the other in Elkton, Virginia. The Miller breweries are dispersed across the United States; however, most markets in the data set do not include a brewery within them. There is one brewery that existed during the study period in Tennessee but closed prior to the merger in 2006. There is dispersion in the size of these geographic markets, with some containing multiple counties and large population centers within them.

1.3.3 Summary Statistics

I provide the summary statistics for each dataset here for the overall market, for Miller, for Coors and for the combined company MillerCoors for 2002–2012. Of these statistics, I emphasize the measures of distance of the nearest brewery to the nearest market, which the company argued was the main impetus for the merger, and product variety. I show the first

Figure 1.3: Location of Miller and Coors Breweries

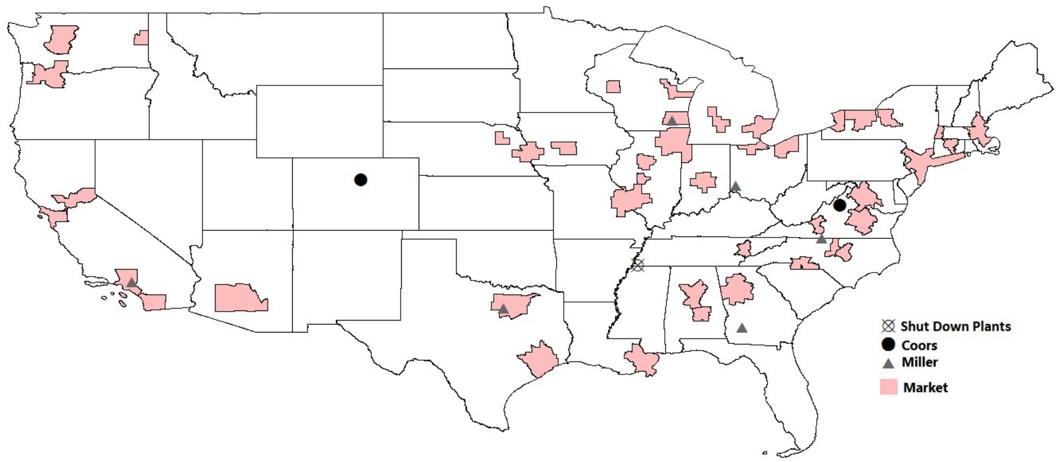

Note: This map provides the locations of Miller breweries, denoted by the gray triangles; the location of Coors breweries, denoted by the black circles; and the markets, denoted by the outlines. Adjacent markets are combined within the data. The Memphis, Tennessee, Coors plant was shut down in 2006, prior to the merger, and is denoted by a crossed-out circle.

main results of the paper, the raw change in product and brand variety, and show that

I first define the difference between product variety and brand variety and examine basic trends to see how the market changed before and after the merger for Miller and Coors. Products in the dataset are defined as a brand × size × packaging type, while brands are names given to products given in the dataset. I consider both definitions for three reasons. First, brands are easily identifiable and clear distinguished in the dataset. For example, Keystone is a different brand from Keystone Light (a lower-calorie version), which is a different brand from Keystone Ice (a version with a higher alcohol by volume). Second, the brand is the highest level of product identification in the dataset beneath the product vendor. If consumers have strong preferences over packaging, such as for twelve packs over twenty-four packs, these results would provide an upper bound on consumer impacts and changes in product variety. Finally, this observable heterogeneity allows me to differentiate between product lines and specific products, which may be important for litigators.

Table 1.1 shows summary statistics for key variables such as revenue, concentration mea-

Table 1.1: Summary Statistics

Variables	Mean	Standard Deviation	Min	Max
Price of good per product	8.6431	4.6013	0.01	282.5500
Parent company national market share	0.1726	0.1474	0.00	0.3977
Industry HHI	2085.7500	146.2158	1864.01	2304.3722
Parent company regional market share	0.1848	0.1678	0.0000	0.6402
Industry regional HHI	2438.8231	707.3270	1082.2442	4387.0833
Number of Miller products, national	92.6220	27.1676	25.0000	177.0000
Number of Coors products, national	48.2464	12.8936	22.0000	97.0000
Number of MillerCoors products, national	161.8263	35.4402	61.0000	254.0000
Number of Miller brands, national	26.6053	6.6754	10.0000	42.0000
Number of Coors brands, national	12.9904	3.3307	6.0000	23.0000
Number of MillerCoors brands, national	48.9947	8.8273	22.0000	70.0000
Minimum distance from a Miller or Coors brewery, in miles	271.2110	218.0582	12.5304	949.1906
Distance from a Miller brewery, in miles	298.5952	238.1581	12.5304	949.1906
Distance from a Coors brewery, in miles	538.0365	245.8177	77.6465	990.6055
Change in distance from a Miller brewery after merger, in Miles	27.3842	51.1858	0.0000	138.7045
Change in distance from a Coors brewery after merger, in miles	266.8255	244.9850	0.0000	796.9849

Note: This table provides summary statistics for the IRI dataset for Miller and Coors beers and the associated companies between 2002 and 2012. The change in distance is calculated based on the linear distance between a brewery and the centroid of the designated market. The average price is calculated over all beers in every market considered in the study.

sures and prices. Of key importance for this table are the measures of the number of Coors brands, Miller brands and distance to the breweries for each. Miller had more breweries, leading to an average distance from a Miller brewery to a market of 314 miles. Potentially due in part to this increased capacity, the number of Miller brands is much greater than the number of Coors brands. On average, Miller supplies 27 brands to the entire country, while Coors supplies only 13 brands. The lack of capacity and distance from markets could potentially affect the cost of producing new products from these sites.

1.3.4 *Merger's Effect on National Brand Variety*

I next compare how Miller and Coors brand variety changed after the merger. In Figure 1.4, I examine the total number of unique brands offered each year for Miller, Coors and MillerCoors. I designate whether a brand is a Miller or Coors brand based on the parent company of the brand prior to the merger. I find that, while brand variety for Miller and Coors declined after the merger, the overall negative effect of the merger was mitigated by new product introductions by the combined company.

Figure 1.4: Number of Unique Brands

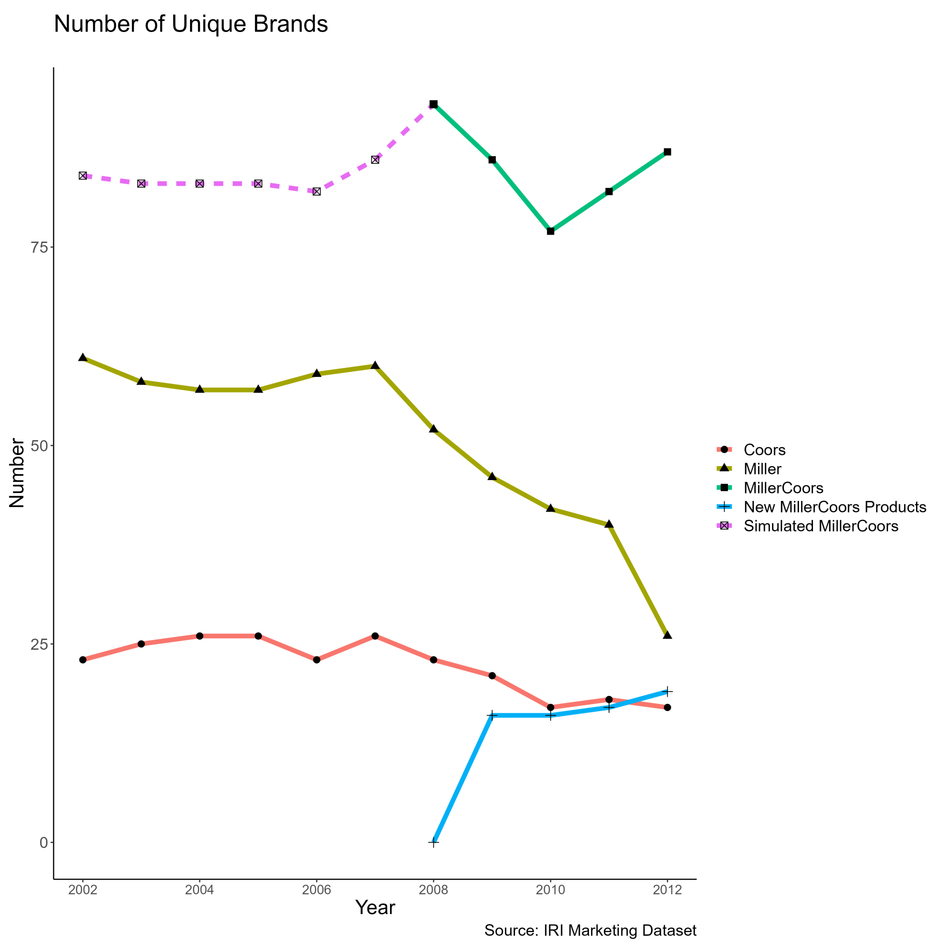

Note: This figure depicts the total number of unique brands offered each year by Miller, Coors and MillerCoors. "MillerCoors" denotes all brands within the new merged company, starting in 2008, and "Simulated MillerCoors" denotes the sum of "Miller" and "Coors" brands prior to 2008. "New MillerCoors products" denotes all new brands created by MillerCoors. This figure depicts all beer brands.

I can further examine the effects of the merger by examining the change in new brands over time. Figure 1.5 depicts this change. Here, new brand offerings declined in the first few years after the merger before increasing again. Prior to the merger, brand offerings remained stable until 2007, when there was a slight increase. Starting in 2009, fewer new brands are created. This effect disappeared by 2011 and 2012, with a large increase in the number of new brands.

Figure 1.5: Changes in Brand Offerings Over Time

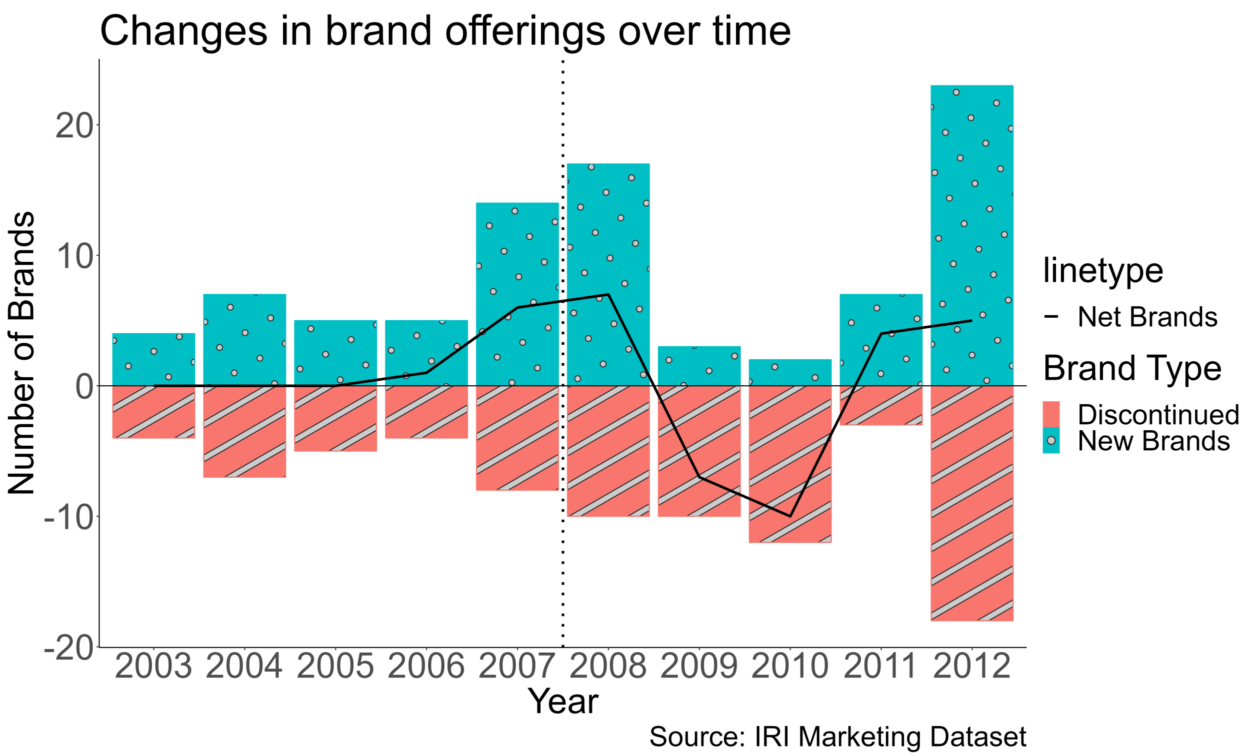

Note: This graph shows the change in number of brands from the prior year. For years prior to 2009, "new brands" and "discontinued brands" are the sum of the Coors and Miller new brands and discontinued brands, respectively. For years from 2009 onward, "new brands" and "discontinued brands" are for MillerCoors only. The figure for 2008 is calculated as the difference in number of brands of Miller and Coors separately and the number in new merged company. The average total number of Miller and Coors brands prior to the merger is 39, and the average number of MillerCoors brands after the merger is 48.

The decline in the number of new brands does not appear to be driven by a small group of markets, as shown by Figure 1.6. The histograms suggest a decline in brand offerings across all markets after the merger. Prior to the merger, for the entire market, with the exception of one market, Coors offered at least 10 brands while offering anywhere from 4 to 11 brands in the twelve-pack market. For Miller, for the entire market, there were 17–27 brands offered. For Coors brands, after the merger, most cities had between 7 and 12 brands for the entire market. For Miller brands, after the merger, most cities had between 15 and 20 brands for the entire market. The decline seems to be stronger for Miller brands, as noted before.

Figure 1.6: Number of Coors and Miller Brands Per Market

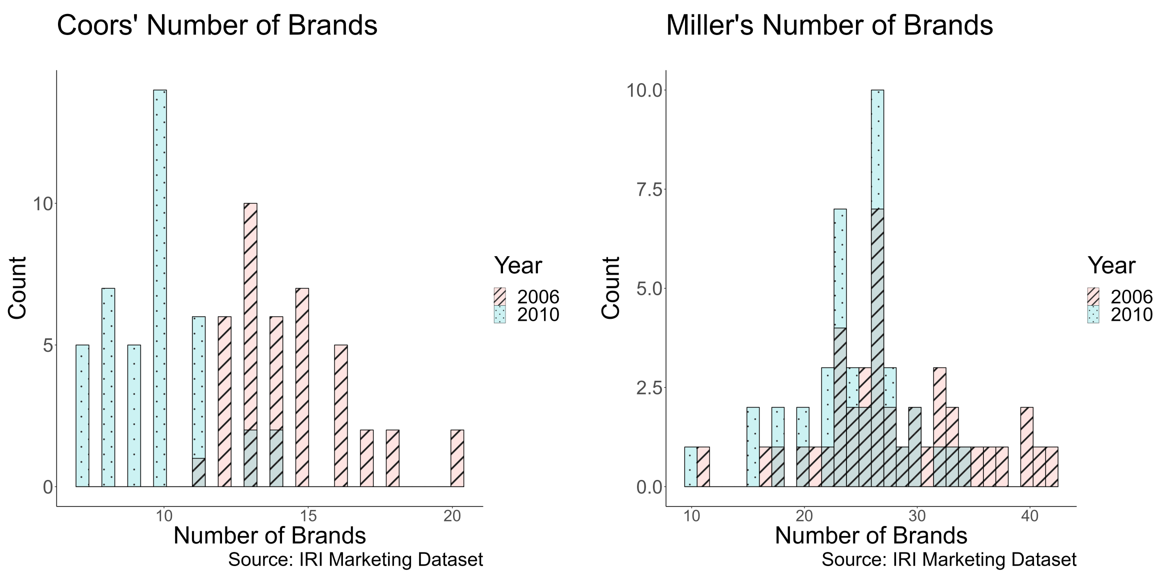

a) b)

Note: This histogram shows the number of brands sold in each market that are designated as sold by Miller or Coors (Miller or Coors brands, respectively) in 2006 and 2010. This does not include brands created by MillerCoors after the date of the merger.

Figure 1.7 shows the average new brands over markets and shows no increase until 2012. Here, it appears that the merger had no immediate effect or, at best, had a lagged effect that did not occur until the end of the sample.

These results show that there were some brand variety changes after the merger. In the context of the basic intuition, it is possible that some brands were dropped due to high costs or increased similarities with other existing brands while brands were introduced to capture new markets or had sufficiently low costs. Since there is enough heterogeneity in the brand changes, it is possible to estimate what these baseline effects are.

Figure 1.7: Change in Brand Offerings Over Time, Market Means

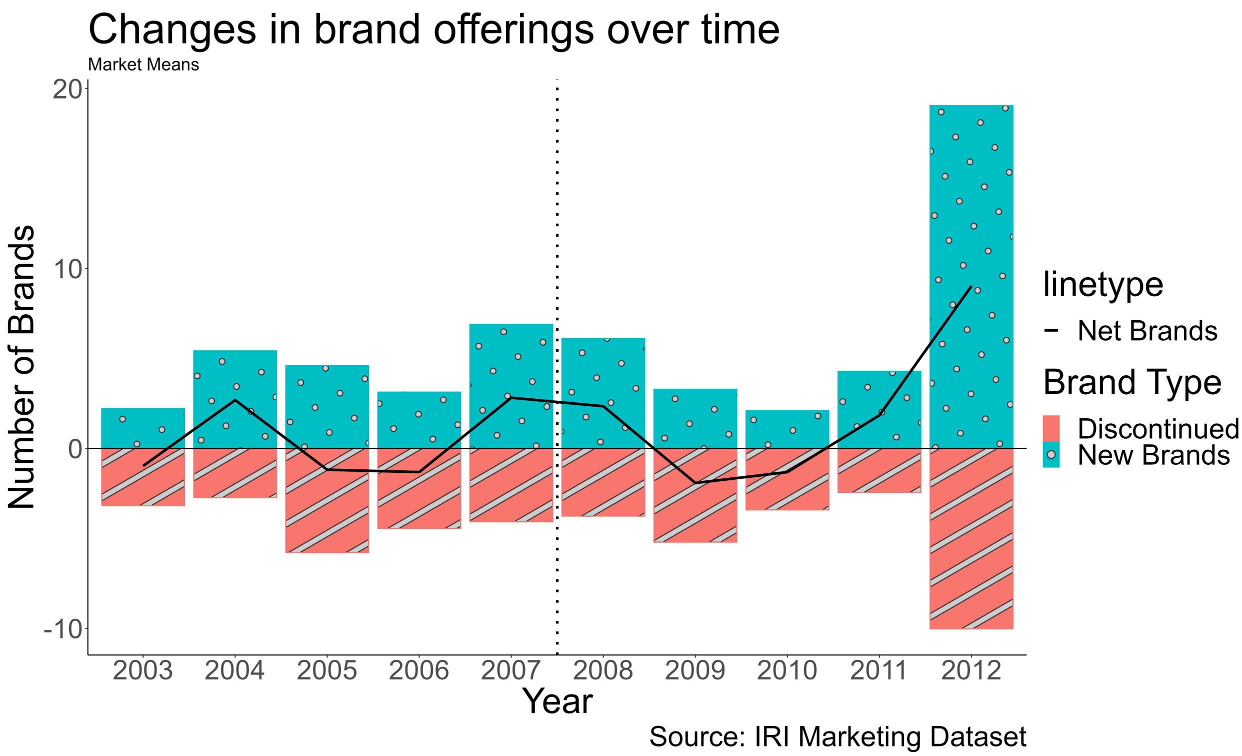

Note: This graph shows the change in brands from the prior year. Prior to 2009, "new brands" and "discontinued brands" are the sum of the Coors and Miller new and discontinued brands, respectively. From 2009 onward, "new brands" and "discontinued brands" are for MillerCoors only. The figure for 2008 is calculated as the difference in the number of brands for Miller and Coors separately and the number for the new merged company. The average total number of Miller and Coors brands prior to the merger is 39, and the average number of MillerCoors brands after the merger is 48.

1.4 Reduced-Form Model

While the summary statistics suggest that the merger's effect on brand variety is positive across all markets, it is less clear whether any particular underlying time- or market-level trends affect these results and what these changes imply for products. We first control for potential confounding factors regarding the relationship between the merger and brand variety. I explain this basic linear model, justify the assumptions and show the results of the model: namely, a decline in brand variety but a null effect on product variety, implying brand consolidation and production of new products in existing brands. I next describe the second model used to control for potential confounding factors regarding the relationship between the merger and brand variety: a difference-in-differences model akin to the model in Ashenfelter, Hosken, and Weinberg (2015). I then explain this model, justify the assumptions and show the results: specifically, a large decline in brand variety relative to that of other top competitors. I finally correct for some issues with these models in several robustness checks.

1.4.1 Model Description

Although the changes in brand variety described in previous sections suggest some changes in brands after the merger, this could be related to changes over time or specific market characteristics. To better control for this possibility, I estimate the following linear model by market m and month t:

$$\log(\text{Number of Brands}_{mt}) = \alpha_{mt} + \beta_1(\text{Post Merger}_t) + \epsilon_{mt} \tag{1.1}$$

where α_{mt} is state–time fixed effects, Post Merger_t is a dummy for whether the observation falls at or after July 1, 2008 (denoted the postmerger period), and ϵ_{mt} is a term representing unobserved market-specific heterogeneity at the market–month–year level. This basic model tests what percentage change in the number of brands is affected by the merger, controlling for observed heterogeneity. Another variant of the model that I also test includes separate dummies on each year to see whether there are any notable changes per year and to better match the pattern seen in the summary statistics.

To precisely estimate the coefficients of the model, I implicitly assume no changes to the underlying market happening before and after the merger. Examining brand variety prior to the merger provides evidence justifying this assumption. Figure 1.8 shows the average brand variety prior to the merger in 2008. Up until the start of 2008, the graph is mostly flat, except for some increases in 2004 and prior to the merger. Even then, any changes are fairly small—of approximately 3–5 brands at most.

Figure 1.8: Average Monthly Brand Variety Prior to the Merger, Miller and Coors

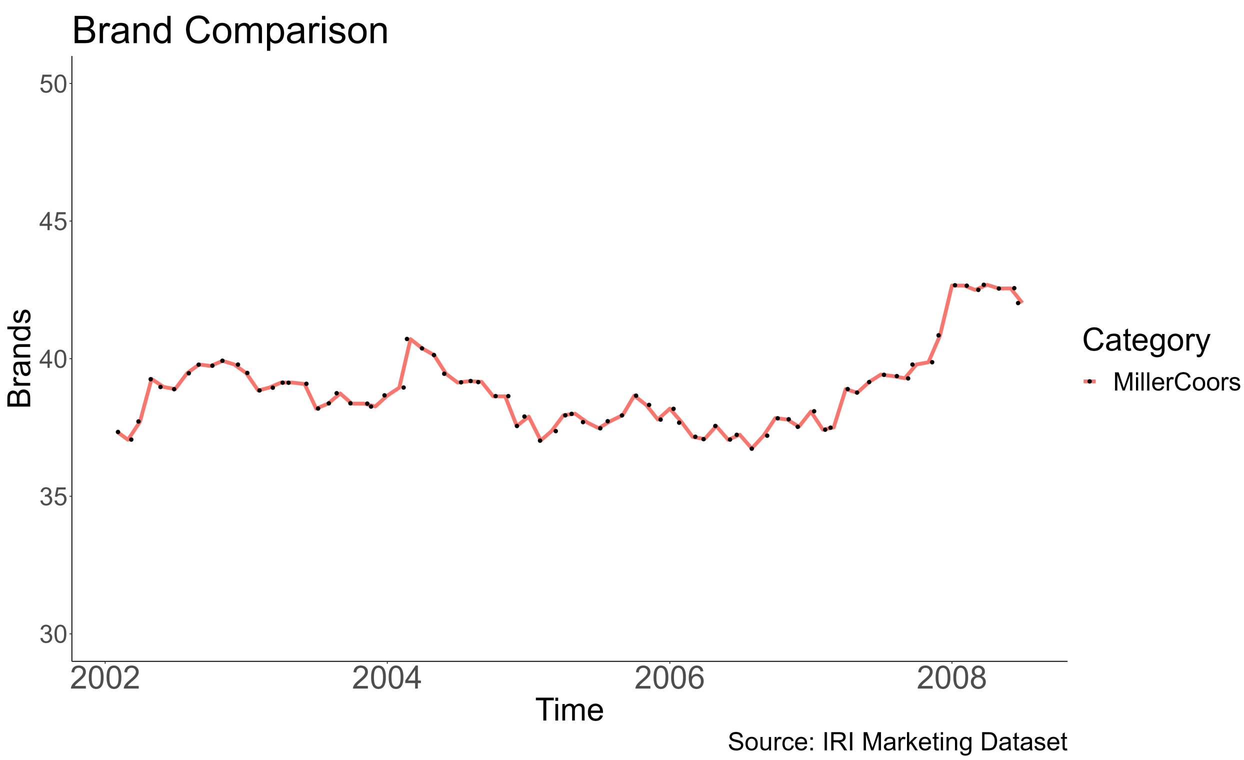

Note: This graph measures the average monthly brand variety for Miller and Coors from January 2002 to May 2007. Brand variety is measured by the number of unique brands that each company has.

1.4.2 Linear Model Results

I first estimate the effect of the merger on brand variety before and after the merger. Table 1.2 shows the results of the linear model.

Table 1.2: Linear Model Results, Product Variety as Number of Brands

	Dependent variable:				
	log(Number of Brands)				
	(1)	(2)	(3)	(4)	(5)
Postmerger	0.097***	0.096***	0.096***	−0.009	−0.022***
	(0.005)	(0.005)	(0.003)	(0.008)	(0.008)
Constant	3.647***	3.649***	3.568***	3.571***	3.567***
	(0.003)	(0.009)	(0.009)	(0.008)	(0.009)
State FE	No	No	Yes	Yes	Yes
Month FE	No	Yes	Yes	No	Yes
Year FE	No	No	No	Yes	Yes
Observations	4,978	4,978	4,978	4,978	4,978
R^2	0.069	0.071	0.706	0.741	0.744
Adjusted R^2	0.069	0.068	0.704	0.739	0.742
Note:	*p<0.1; **p<0.05; ***p<0.01				

Note: Observations are at the market–month–year level."Postmerger" indicates the period after July 2008, the start of the merged company's operation. "Number of Brands" indicates the total number of brands produced by Miller and Coors from 2002 to 2012. These results include all nonstatewide markets and all brands produced by Miller and Coors. HC1 robust standard errors used.

These results are mixed but ultimately show evidence against the positive effect found in the summary statistics. The effect decreases from a gain of 9.7% to a loss of 2.2%, implying that market and time effects play a large role in explaining the brand variety changes after the merger. To put this latter term in perspective with a back-of-the-envelope calculation, after the merger, markets lost on average 2.2% of brands postmerger or approximately one brand.

I next test whether these results hold under a finer definition of products. Rather than use brands, I use the data on products, which combine brand, product packaging, and ounce-size information. The implication of these results is that, while brands may have changed, there may have been an increase in other products. Table 1.3 shows the model under this finer definition.

Table 1.3: Linear Model Results, Product Variety as Number of Products

	Dependent variable:				
	log(Number of Products)				
	(1)	(2)	(3)	(4)	(5)
Postmerger	0.067***	0.067***	0.067***	0.017*	0.011
	(0.007)	(0.007)	(0.003)	(0.009)	(0.010)
Constant	4.820***	4.819***	4.764***	4.780***	4.777***
	(0.004)	(0.012)	(0.008)	(0.008)	(0.010)
State FE	No	No	Yes	Yes	Yes
Month FE	No	Yes	Yes	No	Yes
Year FE	No	No	No	Yes	Yes
Observations	4,978	4,978	4,978	4,978	4,978
R^2	0.021	0.022	0.738	0.745	0.746
Adjusted R^2	0.021	0.020	0.736	0.743	0.744

Note: *p<0.1; **p<0.05; ***p<0.01

Note: Observations are at the market–month–year level."Postmerger" indicates the period after July 2008, the start of the merged company's operation. "Number of Products" indicates the total number of brands produced by Miller and Coors from 2002 to 2012. These results include all nonstatewide markets and all products. HC1 robust standard errors used.

The ultimately null results of the merger's effect on product variety combined with the results on brand variety above provide some evidence of product consolidation. With no controls, the effect is positive and implies a 5% increase in product variety due to the merger.

However, with the introduction of year fixed effects, the point estimate of the effect of the merger on product variety is 0.1% and indistinguishable from zero. Taken with the brand results, this implies that the firm may have dropped a brand to create more products within its main brands. These results indicate that the decline in brands was offset by an increase in products, leading to the ultimately null result.

Overall, these results imply a slight decline in the number of brands offered by Miller and Coors offset by potential increases in products. However, there are shortcomings of this basic model. First, there are no controls for competitive effects. There could be industry-specific effects from the merger that influence the results, and a comparison examining only the firms involved ignores this possibility. Second, there could be other changes happening around the merger that the linear model would not capture well. For example, the great recession occurred during the sample period. If it had an effect on demand for alcohol during this time compared to demand prior to the merger, this could bias the results upwards. Therefore, another model is needed to better control for these industry-wide effects.

1.4.3 Difference-in-Differences Model

I describe the difference-in-differences model here, which uses a subset of the firm's data alongside the Miller and Coors data to estimate how MillerCoors brand offerings compare with the offerings of competitors. This model is used for comparing effects on brand variety relative to that of competitors, which may be important for policymakers. Here, I describe the variables, estimation strategy and groups that I use for comparison.

I use the model below to estimate the impact of the merger on the number of brands of each firm i in each market m at each period t:

$$\log(\text{num brands})_{imt} = \beta_1(\text{Post Merger}_t) + \beta_2(\text{Miller}_i + \text{Coors}_i) + \beta_3(\text{MillerCoors}_i) + \epsilon_{imt} \quad (1.2)$$

where α_{imt} represents market, firm and time fixed effects, (Post Merger_t) is an indicator for whether the observation is after the completion of the merger, $\text{Miller}_i + \text{Coors}_i$ is a sum of the indicators for the Miller and Coors brands, and MillerCoors_i is an indicator for whether the brand is a MillerCoors brand. The last coefficient acts as a difference-in-differences coefficient of interest, as it measures the additional impact of the number of brands after the merger and under the merged company, relative to the change in the control group.

I employ two different control groups to compare the effects of the merger on brand variety. One important distinction is that firms may act differently in terms of brand variety depending on their size. Craft breweries, which are typically small in overall market share, have only so much capacity to not only provide variety but ship it to every market in the United States. Therefore, such breweries are typically local. To deal with this issue, I examine two groups: all firms in the dataset and firms in the top 5 quintile of market share, which is the set of firms used in Ashenfelter, Hosken, and Weinberg (2015).

The difference-in-differences model implicitly assumes that both Miller and Coors and their competitors were similar prior to the merger. I provide evidence that the difference-in-differences assumption is valid when I use the top 5% of firms in the market. I first examine how MillerCoors compares with all firms. However, this may not be an appropriate comparison. Many firms, such as craft brewing companies, are small and produce only a few brands, such as regional brands. First, I plot the average MillerCoors brand variety across markets versus the average brand variety across all firms and all markets. Figure 1.9 displays the comparison of Miller and Coors brand variety versus that of all firms.

As seen in this figure, this comparison group greatly differs in levels and may not be an appropriate group for comparison. Much of this variation comes from smaller firms in the dataset, such as craft breweries.

I next examine how MillerCoors compares with firms in the top 5% market share quartile, which is analogous to the comparison group used in Ashenfelter, Hosken, and Weinberg

Figure 1.9: Comparison of Brands Offered by MillerCoors and All Breweries in IRI Database

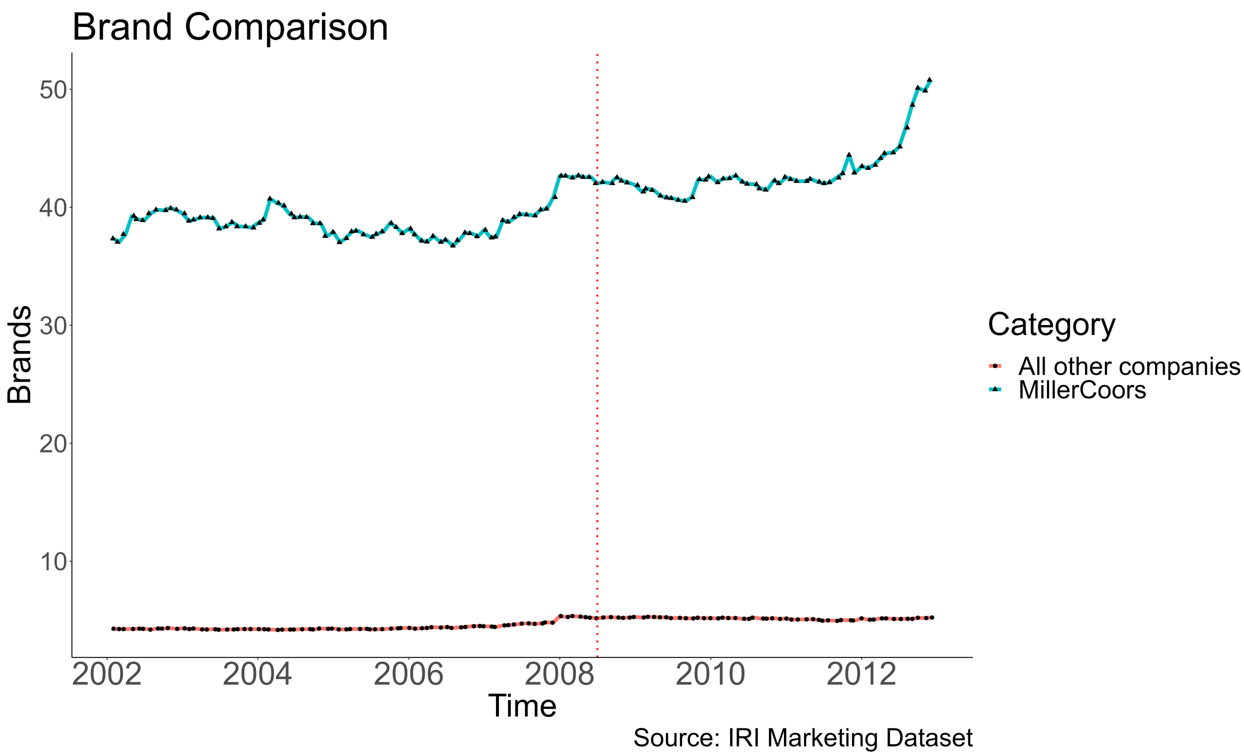

Note: Observations are at the month–year level. This graph plots the average of the number of brands offered over all markets. Miller and Coors are combined as one company prior to their merger. The red dashed line denotes the date that the MillerCoors merger was finalized.

(2015). I focus on the top 5% as it may be a better comparison group for Miller and Coors, which were the second and third largest breweries in the country prior to the merger. This comparison group has been used in other works in the literature (Miller and Weinberg (2017)). First, I plot the average MillerCoors brand variety across markets versus the average brand variety across all firms in the top 5% market share quartile and all markets. Figure 1.10 displays the comparison of Miller and Coors brand variety versus that of all firms with market shares in the top 5

Figure 1.10: Comparison of Brands Offered by MillerCoors and All Breweries Within Top 5% of Market Share

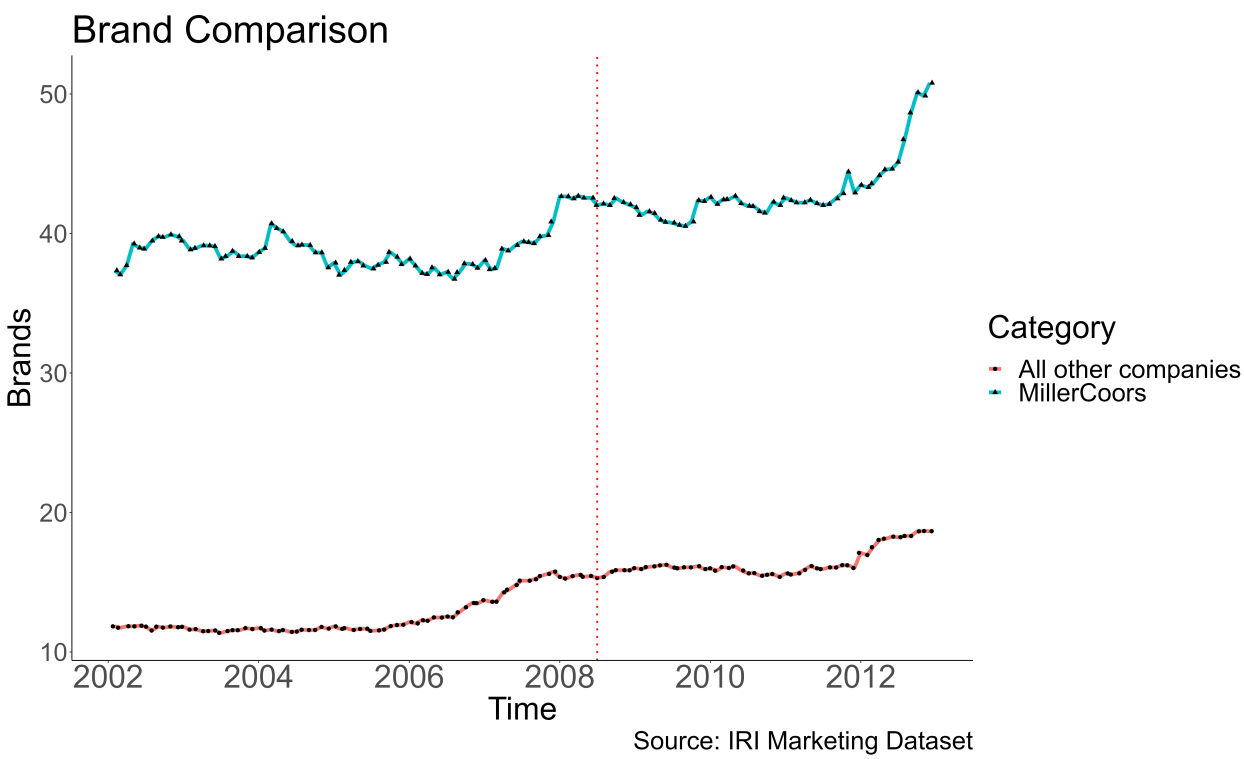

Note: Observations are at the month–year level. This graph plots the average of the number of brands offered over all markets. Miller and Coors are combined as one company prior to their merger. The red dashed line denotes the date that the MillerCoors merger was finalized.

With the exception of a premerger increase driven by rapid growth by Anheuser-Busch in 2007, the graphs are far closer in levels than the previous figures. While the inclusion of Anheuser-Busch is a concern given its rapid growth prior to the merger, removing it from the group does not change the sign of the results below.

1.4.4 Difference-in-Differences Model Results

The impact of the merger in a comparison with the outcomes in the first group may not represent the true outcome, as seen in Table 1.4.

Table 1.4: Difference-in-Differences Results, All Breweries in IRI Database

	Dependent variable:				
	log(Number of brands)				
	(1)	(2)	(3)	(4)	(5)
Postmerger	1.078***	0.078***	0.073***	−0.011	−0.012
	(0.017)	(0.016)	(0.016)	(0.010)	(0.010)
Miller or Coors	2.646***	2.646***	2.646***	2.648***	2.648***
	(0.025)	(0.025)	(0.026)	(0.026)	(0.026)
MillerCoors	0.020	0.019	0.025	0.022	0.022
	(0.018)	(0.018)	(0.018)	(0.017)	(0.017)
State FE	No	No	Yes	Yes	Yes
Month FE	No	Yes	Yes	No	Yes
Year FE	No	No	No	Yes	Yes
Observations	298,153	298,153	298,153	298,153	298,153
Adjusted R^2	0.605	0.605	0.606	0.607	0.607

Note: *p<0.1; **p<0.05; ***p<0.01

Note: Observations are at the firm–market–month–year level. "Postmerger" indicates the period after July 2008, the start of the merged company's operation. Standard errors are clustered at the market level. For the last set of regressions, to remove colinearity with the year, month and postmerger variable, the December fixed effect is removed.

Similarly to the results on brand variety change from the linear model, there is no significant effect of the merger on brand variety, and much of the effect is on the firm itself. The point estimates for the difference-in-difference coefficient, MillerCoors, is 0.02, which would imply a 2% change in the number of brands offered after the merger for MillerCoors relative to the number for all competitors in the market. However, this coefficient is not significant, regardless of inclusion of state, month and year fixed effects.

However, the effect of the merger on brand variety differs when we directly compare the MillerCoors brand variety changes with the outcomes of other top competitors in the beer industry. Table 1.5 shows the results from the difference-in-differences model with the comparison to the breweries within the top 5% of market share:

Table 1.5: Difference-in-Differences Results, All Breweries Within Top 5% of Market Share

	Dependent variable:				
	log(Number of brands)				
	(1)	(2)	(3)	(4)	(5)
Postmerger	2.485***	0.319***	0.318***	0.030***	0.018***
	(0.023)	(0.019)	(0.019)	(0.006)	(0.006)
Miller or Coors	1.479***	1.479***	1.479***	1.481***	1.481***
	(0.027)	(0.027)	(0.027)	(0.027)	(0.027)
MillerCoors	−0.221***	−0.221***	−0.220***	−0.222***	−0.222***
	(0.024)	(0.024)	(0.024)	(0.024)	(0.024)
State FE	No	No	Yes	Yes	Yes
Month FE	No	Yes	Yes	No	Yes
Year FE	No	No	No	Yes	Yes
Observations	45,712	45,712	45,712	45,712	45,712
Adjusted R^2	0.916	0.916	0.917	0.918	0.918

Note: *p<0.1; **p<0.05; ***p<0.01

Note: Observations are at the firm–market–month–year level. "Postmerger" indicates the period after July 2008, the start of the merged company's operation. Standard errors are clustered at the market level. For the last set of regressions, to remove colinearity with the year, month and postmerger variable, the December fixed effect is removed.

Here, I find a large and significant decline in brand variety for MillerCoors compared to that of its competitors. The point estimates for the difference-in-difference coefficient, MillerCoors, is -0.22, which would imply a 22% decline in the number of brands offered after the merger for MillerCoors compared to the number for all other competitors in the

AHW group. This effect remains significant with the inclusion of year, month and state fixed effects.

Finally, to verify how these results can compare against the linear model, I examine the impact over time for just the top quartile results. This is to see whether the positive effect is driven by brand increases in the later years or at the start of the postmerger period. I use the model below, with observations for firm i at time t in market m:

$$\begin{aligned}\log(\text{Number of Brands}_{imt}) = \alpha_{mt} + \beta_0(\text{Miller}_i + \text{Coors}_i)+ \\ (\text{Year\_1\_After\_Merger}_t)\beta_1 + (\text{Year\_1\_After\_Merger}_t * \text{MillerCoors}_i)\beta_2+ \\ (\text{Year\_2\_After\_Merger}_t)\beta_3 + (\text{Year\_2\_After\_Merger}_t * \text{MillerCoors}_i)\beta_4+ \\ (\text{Year\_3\_After\_Merger}_t)\beta_5 + (\text{Year\_3\_After\_Merger}_t * \text{MillerCoors}_i)\beta_6+ \\ (\text{Year\_4\_and\_Above\_After\_Merger}_t)\beta_7+ \\ (\text{Year\_4\_and\_Above\_After\_Merger}_t * \text{MillerCoors}_i)\beta_8 + \epsilon_{imt}\end{aligned} \tag{1.3}$$

The main difference between this model and the previous one is that, here, I interact each year after the merger with the MillerCoors dummy variable. After the fourth year, I group all observations together. The results for this model adding month and state fixed effects are in Table 1.6.

These results grow increasingly negative over time. Under month and year fixed effects, the difference-in-differences coefficient interacted by years after the merger grows from -0.20 to -0.23. This implies that the effect of the merger grew over time, reaching its largest toward the end of the sample period. In percentage terms and as a back-of-the-envelope calculation using the mean number of MillerCoors brands per market, these coefficients imply a loss of 10–11 brands relative to the change in brand number of competitors.

Table 1.6: Difference-in-Differences Results, Effect Separated by Year

	Dependent variable:			
	log(Number of brands)			
	(1)	(2)	(3)	(4)
Year 1 $*$MillerCoors	-2.377^{***}	-0.209^{***}	-0.209^{***}	-0.209^{***}
	(0.020)	(0.020)	(0.020)	(0.020)
Year 2 $*$MillerCoors	-2.393^{***}	-0.224^{***}	-0.225^{***}	-0.224^{***}
	(0.019)	(0.024)	(0.024)	(0.024)
Year 3 $*$MillerCoors	-2.383^{***}	-0.214^{***}	-0.216^{***}	-0.214^{***}
	(0.020)	(0.026)	(0.026)	(0.026)
Year 4+ $*$MillerCoors	-2.397^{***}	-0.227^{***}	-0.229^{***}	-0.227^{***}
	(0.025)	(0.031)	(0.031)	(0.031)
State FE	No	Yes	No	Yes
Month FE	No	No	Yes	Yes
Observations	45,712	45,712	45,712	45,712
Adjusted R^2	0.575	0.917	0.916	0.917
Note:		*p<0.1; **p<0.05; ***p<0.01		

Note: Observations are at the firm–market–month–year level. The year variables indicate the period after July 2008, the start of the merged company's operation. Only the ten largest firms by national revenue shares prior to 2007 are included in this regression. Standard errors are clustered at the market level. For the last set of regressions, to remove colinearity with the year, month and postmerger variable, the December fixed effect is removed.

1.4.5 Variants on the Control Group

While the results above show how the number of brands compares against that of firms not part of the merger (the control group), there are concerns about the control group used and whether the results would still hold under corrections to the control group. In this section, I explain these issues and the two techniques that I use to deal with them.

1.4.6 Key issues with the merger and comparisons

There are two main issues with the comparison group, which the following two techniques attempt to deal with. Both issues are related to the parallel trends assumption, specifically, how the control group compares in levels and in pretrends. First, most companies produce fewer brands than Miller, Coors and Anheuser-Busch, the top three firms in market share. Figure 1.11 shows the average number of brands for the top 5 firms in market share during the 2002–2012 period. Given this issue, comparisons in levels may be difficult. While reducing the sample to firms in the top 5 percentile of market share does alleviate this issue and allow better comparisons between large, similar firms, this leads to a second issue regarding parallel trends.

The second issue is the differential pretrend that we observe for the control group for years prior to 2008. The majority of the increase prior to the merger is attributable to an increase in Anheuser-Busch's brand variety in 2006. Additionally, Anheuser-Busch and InBev merged in 2007; however, InBev is a Belgian company and had no breweries located in the United States prior to the merger.[14] While removing these companies from the data could help, Anheuser-Busch InBev is the top competitor and has the largest number of brands, and this raises the issues with using comparable levels and companies described previously. Therefore, removing Anheuser-Busch and Inbev is not desirable. In the following sections

[14]Since the focus of this chapter is on a large domestic merger with cost synergies, this merger is not discussed here. See *NY Times*, "Anheuser-Busch Agrees to Be Sold to InBev".

Figure 1.11: Average Monthly Brand Variety, Top 5 Firms

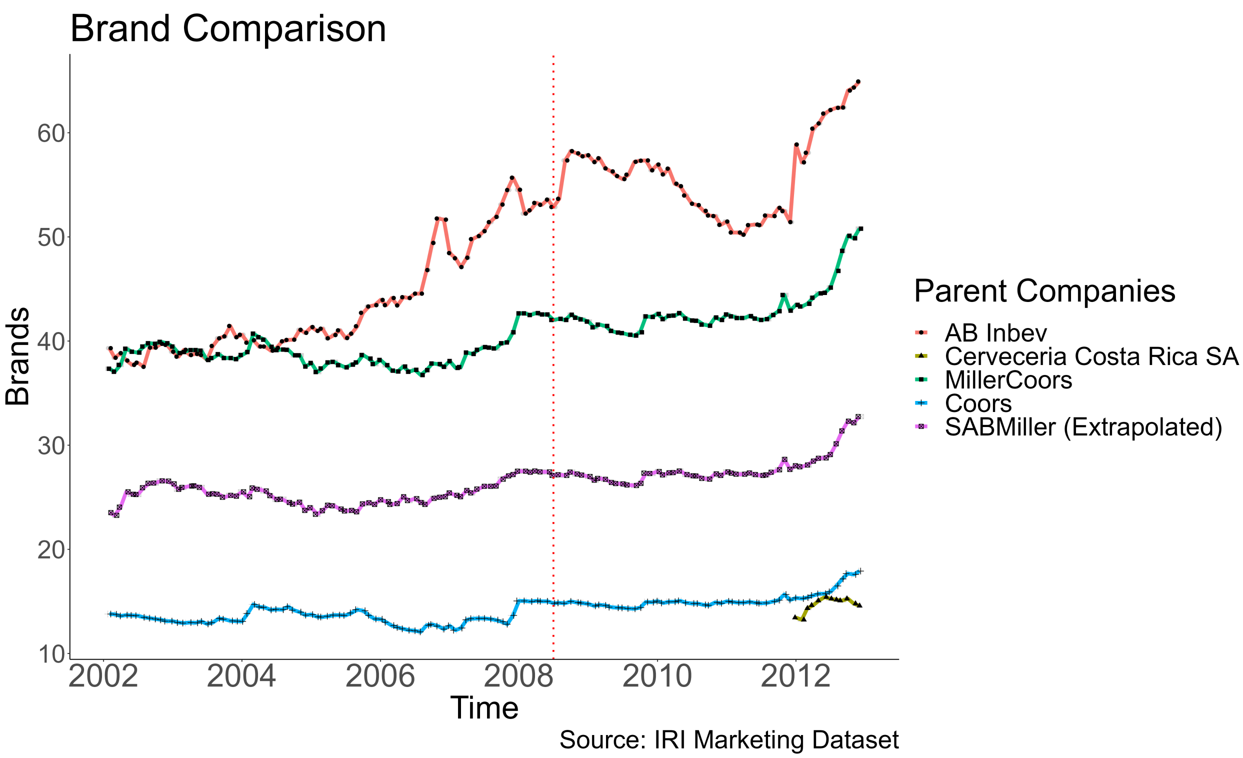

Note: This graph plots the average of the number of brands offered over all markets. Miller and Coors and Anheuser-Busch and InBev are combined as one company prior to their mergers. The Miller and Coors numbers of brands are added to this graph for comparison. Observations are at monthly level. The red dashed line denotes the date when the MillerCoors merger was finalized. The lines in gray are firms in which the minimum number of brands produced at any time is fewer than ten.

and in the appendix, I explain how I correct this by removing the time trend from the data using a strategy from Goodman-Bacon (2021) and synthetic control methods from Abadie and Gardeazabal (2003).

1.4.7 Detrending the data

One method of creating a better control group is to remove the time trend from the data. This method comes from Goodman-Bacon (2021) and involves the following two-step process. In the first stage, I estimate linear trends in the number of brands for the control group and the treatment group using only the pretreatment period. Next, I subtract the predicted time trend for the number of brands from the entire period. This method is designed to remove any trends from the groups that may not be accounted for by the year fixed effects, such as the increase in Anheuser-Busch's brand variety.

Figure 1.12 displays the detrended data for the top 5 percentile comparison group, which does match the pretrend prior to the merger better on levels but fluctuates more in the earlier part of the sample period. Although this method helps match the pretrend better and, based on the figure, MillerCoors still shows a large increase in brand variety after the merger, using these earlier data causes some issue.

Using these detrended data, I estimate a version of the model that accounts for negative brand values that has the detrended number of brands as the dependent variable. The difference-in-difference coefficient, MillerCoors, now represents the number of brands lost because of the merger and because they were part of the MillerCoors company. The results are in Table 1.7:

The effect on the detrended data is negative and larger than the baseline results. After the merger, the number of brands offered by MillerCoors compared to that of its competitors ranges from about -6 to -5.9 and remains significant under standard fixed effects. Net of this time trend, the effect is smaller but still negative.

Figure 1.12: Comparison of Brands Offered by MillerCoors and All Breweries Within Top 5% of Market Share, Detrended

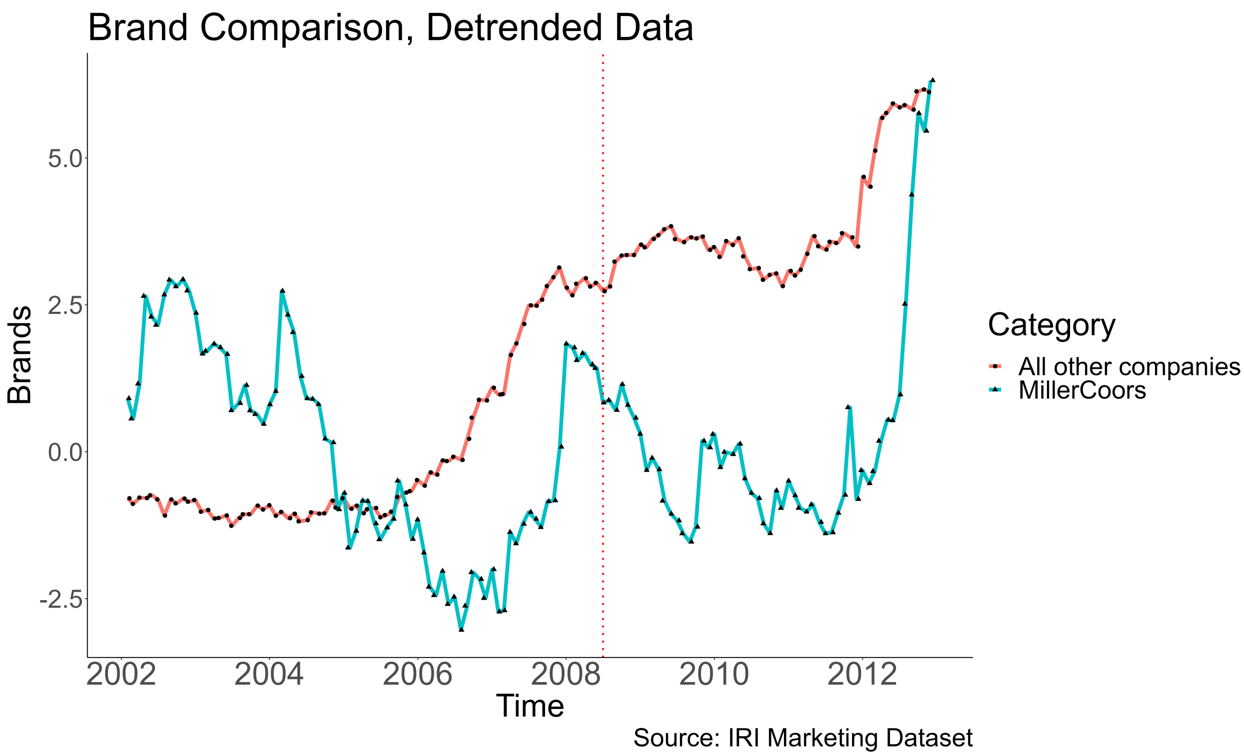

Note: This graph plots the average of the number of brands offered over all markets. Miller and Coors are combined as one company prior to their merger. Miller and Coors are kept separate for comparison. Observations are at the monthly level. The red dashed line denotes the date that the MillerCoors merger was finalized. The detrending procedure is as follows: In the first stage, I estimate linear trends in the number of brands for the control group and the treatment group using only the pretreatment period. I then subtract the predicted time trend from the number of brands for the entire period.

Table 1.7: Difference-in-Differences Results, Using Detrended Data

	Dependent variable:				
	Detrended Number of Brands				
	(1)	(2)	(3)	(4)	(5)
Postmerger	6.758***	6.088***	6.092***	0.922***	0.668***
	(0.372)	(0.252)	(0.252)	(0.114)	(0.118)
Miller or Coors	−0.635	−0.635	−0.632	−0.757	−0.759
	(0.833)	(0.833)	(0.829)	(0.827)	(0.827)
MillerCoors	−5.993***	−5.992***	−5.997***	−5.869***	−5.867***
	(0.557)	(0.557)	(0.555)	(0.556)	(0.556)
State FE	No	No	Yes	Yes	Yes
Month FE	No	Yes	Yes	No	Yes
Year FE	No	No	No	Yes	Yes
Observations	38,749	38,749	38,749	38,749	38,749
Adjusted R^2	0.073	0.073	0.086	0.098	0.098

Note: *p<0.1; **p<0.05; ***p<0.01

Note: Observations are at the firm–market–month–year level. The "Postmerger" variable indicates the period after July 2008, the start of the merged company's operation. Only the ten largest firms by national revenue share prior to 2007 are included in this regression. Standard errors clustered at the market level. For the last set of regressions, to remove colinearity with the year, month and postmerger variable, the December fixed effect is removed.

1.4.8 Summary of the Two Models

Several conclusions can be drawn from these these two classes of models. The linear model shows that Miller and Coors experienced a decrease in brand variety per market after the merger, compared to their brand variety prior to the merger. However, there was no decline in product variety, indicating that some products within established brands were created after the merger to offset the loss of certain brands. This corroborates a finding from prior work that merging firms consolidate product lines in their core competencies.

However, the linear model alone may not be sufficient for understanding overall product variety changes after the merger. I therefore create a difference-in-difference model to control for competitor effects. I find a decrease in brand variety relative to that of top competitors but no decrease in brand variety relative to all brands on the market. Which of these effects seems more plausible is heavily reliant on how the market is defined, a matter on which this chapter takes no stance but provides multiple results to inform the question.

A model is needed to explain how significant this change truly is for consumers. While the data appendix does show that the majority of discontinued brands were low-revenue, low-sales brands, consumer welfare may depend more on the overall products currently offered, and consumers may substitute to existing products in the market. Alternatively, consumers may not care about MillerCoors products, and the gain in product variety from that brand alone may lead to little change in consumer welfare. Therefore, I need a consumer welfare model to see which of these two reduced-form models best aligns with final consumer outcomes. With this in mind, I return to the product level for the structural model.

1.5 Conclusion

In this chapter, I examine how a merger may affect changes in product variety offered to consumers. I first establish the main setting for this chapter: the beer industry and the Miller and Coors merger of 2008. This setting corresponds to a well-established market with a large set of competitors, a merger of two large firms within it, and well-defined product variety that I can easily map to the main data source for this chapter, the IRI marketing database. After providing evidence of national product varieties decreasing and then increasing late in the sample period, I apply two sets of reduced-form models to see how this change compares at the market level for the two companies by themselves and how it compares against similar competitors in the market.

I find that, at the market level, product variety does change after a merger and this change is significant for consumers. I find that, after the merger, while the number of MillerCoors brands per market falls the number of products remains constant, implying a streamlining of product choices at the firm level. However, compared to that of competitors in the market, MillerCoors's brand variety falls by 22%. This negative result is robust to removing the time trend. In the next chapter of the dissertation, I explore what the consumer welfare effects of changes of product variety are in this merger.

1.6 Appendix

1.6.1 Data Cleaning and Summaries

This chapter uses a variety of datasets to fully understand the effect of the MillerCoors merger on product variety. As such, this data appendix contains information on each dataset and what is done to clean the data. Each section describes the dataset, provides summary statistics, and describes the data cleaning process and justification for any changes made. The final section describes how merges between datasets were accomplished.

The IRI Marketing Data Set

The first data source is the IRI Marketing Supermarket data set. The IRI Marketing data set spans from 2001 - 2012 and contains supermarket transaction data from 51 marketing regions. These marketing regions are either singular counties or groups of counties. Each observation is an individual sale of a product with a unique UPC. The sale recorded is the total price of the sale, the number of units of that good purchased, which store (anonymized by year) it was purchased in and which market it was purchased in. Each good is further identified on their brand, vendor, parent company, packaging medium, and the size by ounces.

Due to the size of the data, the data is split into four different groups of files per year. The first group contains the individual sales of each beer as described above. The second group are product information files. There are four of these files: One for 2001 - 2006, one for 2007, one for 2008 - 2011 and one for 2012. For each file, IRI provides product information for each brand within the time frame listed. However, any changes to any brand during this period is overwritten - for example, if a brand were to change parent companies over the course of a year, this would not be observed in the dataset. The third group are the dates files and lists the week that the product was purchased. Using this, I am able to tell which week, month and year the item was purchased in. I use the start of the week to determine what year or month it is in. For years where the first week falls within a different year, I attribute the week to January of the current year. The fourth group is the delivery files, which list the market and anonymized stores where sales were conducted. There is one duplicate store chain, which I remove the earliest version of.

For the purposes of this chapter I remove several markets and a year of data that are problematic for either estimation or data reasons. First, there were several major store level mergers that occurred in 2001. To solely focus on the MillerCoors merger, data from the year 2001 are dropped. To remove potential external constraints on product variety, I do not include markets from states that place restrictions on alcohol varieties (such as limits on ABV) or distribution. These states include Pennsylvania, Kansas, Utah, Oklahoma and Minnesota. I remove Harrisburg, Philadelphia, Tulsa, Oklahoma City, Salt Lake City, and Minneapolis. I remove Providence, Rhode Island since it drops out of the dataset in 2007. Finally, I remove state level markets so census controls on individual markets can be better utilized. This includes Mississippi and South Carolina. This reduces the number of markets in this study to 39 markets in a total of 28 states.

Algorithm to Clean the IRI Dataset

The algorithm to clean the IRI data and prepare it for analysis proceeds as follows: each year, all IRI data for the year is merged together, new variables are generated, then I calculate summary statistics by month and year and output them for later analysis. There are three variants of this algorithm that create summary statistics for the following categories: the entire dataset, for each market and for each firm. Each one uses the following algorithm for merging:

1. I read in the scanner data and drop products that are missing units.
2. I merge the product information to the scanner data.
3. I read in store and market data. I drop stores that are duplicated or have other issues as described in the last section.
4. I drop all markets that have issues, as described in the last section.
5. I combine the data with the dates data to get the exact weeks the products were sold.
6. I attribute any week ending in January to January, and drop data from weeks ending in January of next year.

7. I merge with the file that identifies which products belong to Miller and Coors prior to the merger and which do not.

8. Using the above classification, I identify if goods are from Miller, from Coors, from MillerCoors, and whether they were new MillerCoors goods.

9. I calculate market shares and HHI at the national level and at the market level per year and per month.

10. I calculate summary statistics for the categories described above for month and year. I repeat this process for Miller classified goods only, Coors classified goods only and MillerCoors classified goods only.

11. For certain subsets of the data, I calculate the change in brands and the change in products over time. This is done through getting the list of brands, cleaning both to address any minor text changes (Such as Hamms and Hamm's), and then merging both to find which brands matched (continued), which brands do not match in the latest data (discontinued) and which brands do not match from earlier data (new). I do this for both products and brands. Additional cleaning is done to remove brands that are discontinued twice from the discontinued category and brands that are considered new twice from the new category.

12. The data is saved then used for summary statistics within the paper.

Categorizing Brands in the IRI Dataset

During the period of study, brands may change either their vendor or parent company due to mergers, acquisitions or spinoff subsidiaries. In the results presented in the paper, I use Brand-Parent to categorize goods. For Miller and Coors brands, I use Brand-Parent to categorize goods for each company prior to the merger, and track which goods are still available after the merger. I use Brand-Parent primarily to include subsidiary companies. Subsidiary companies could be a target for the merger, and therefore their product variety may be affected by a merger.

Examining Features of the Discontinued and New Brands

I provide graphs summarizing features of new and discontinued brands below. The features I examine are revenues of new and discontinued brands in their first and last year respectively, the total products under each brand, the total products sold under each brand, the market share of new and discontinued brands in their first and last year respectively, the share that new and discontinued brands contributed to the parent company's market share, the mean number of markets and the mean number of stores.

Figure 1.13: Total Revenue of Discontinued and New Brands

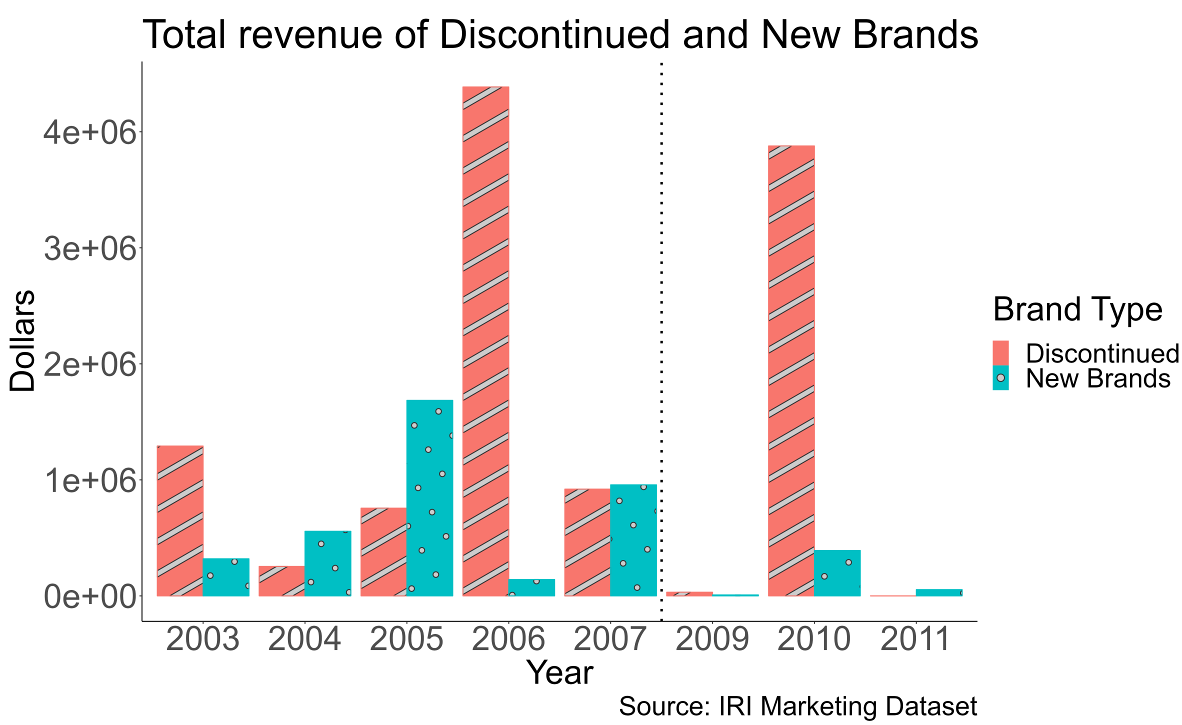

Note: This graph shows information from brands that were added in that year and brands that were discontinued after that year. 2008 is removed due to data issues. Prior to 2009, "new brands" and "discontinued brands" are the sum of the Coors and Miller new and discontinued brands, respectively. From 2009 onward, the "new brands" and "discontinued brands" are for MillerCoors only.

Figure 1.14: Total Products of Discontinued and New Brands

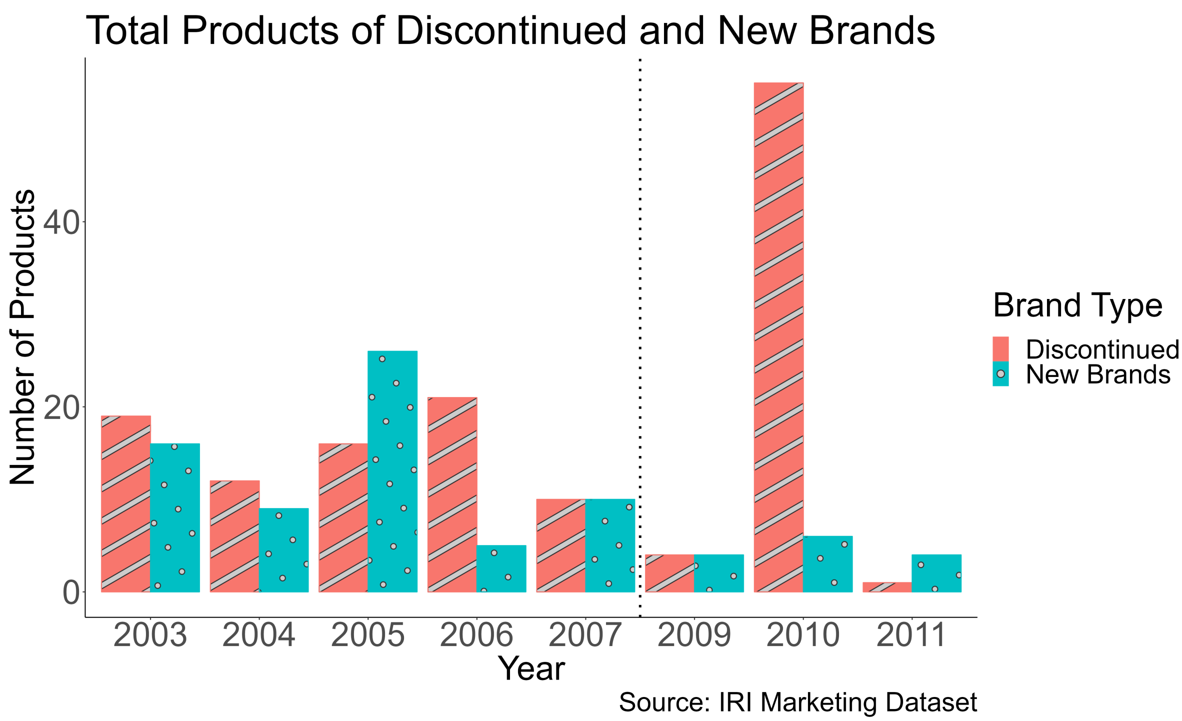

Note: This graph shows information from brands that were added in that year and brands that were discontinued after that year. 2008 is removed due to data issues. Prior to 2009, "new brands" and "discontinued brands" are the sum of the Coors and Miller new and discontinued brands, respectively. From 2009 onward, the "new brands" and "discontinued brands" are for MillerCoors only.

Figure 1.15: Total Sold Products of Discontinued and New Brands

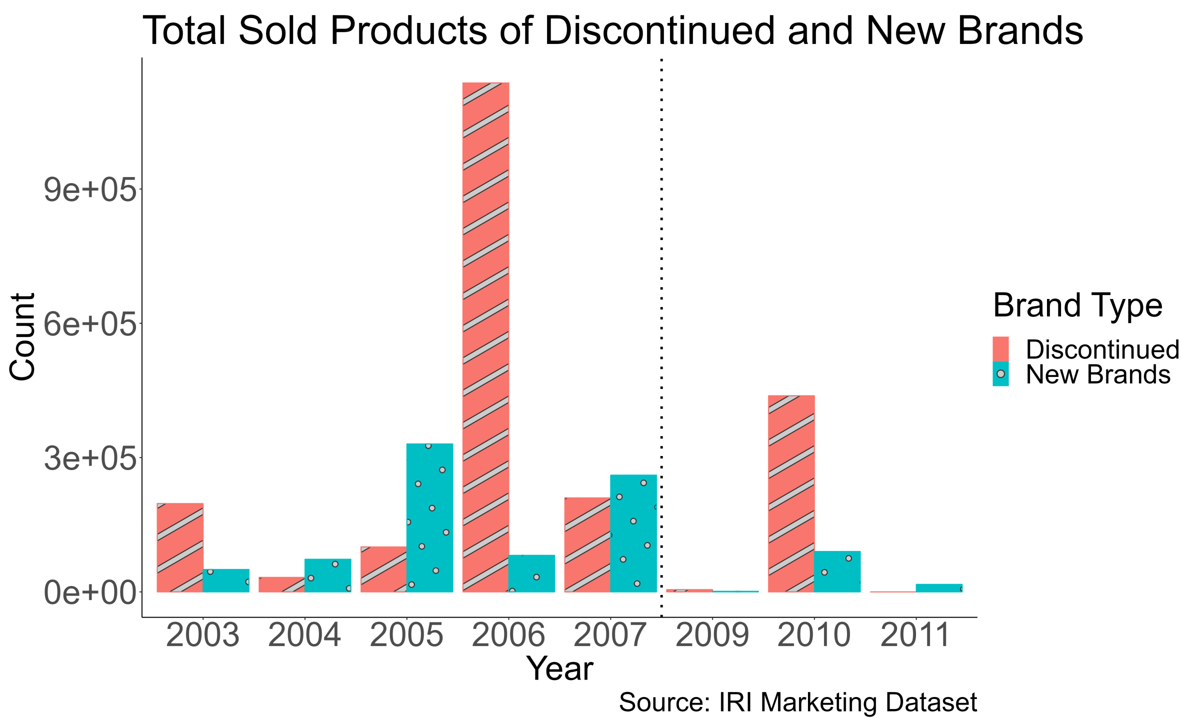

Note: This graph shows information from brands that were added in that year and brands that were discontinued after that year. 2008 is removed due to data issues. Prior to 2009, "new brands" and "discontinued brands" are the sum of the Coors and Miller new and discontinued brands, respectively. From 2009 onward, the "new brands" and "discontinued brands" are for MillerCoors only.

Figure 1.16: Total Market Share Products of Discontinued and New Brands

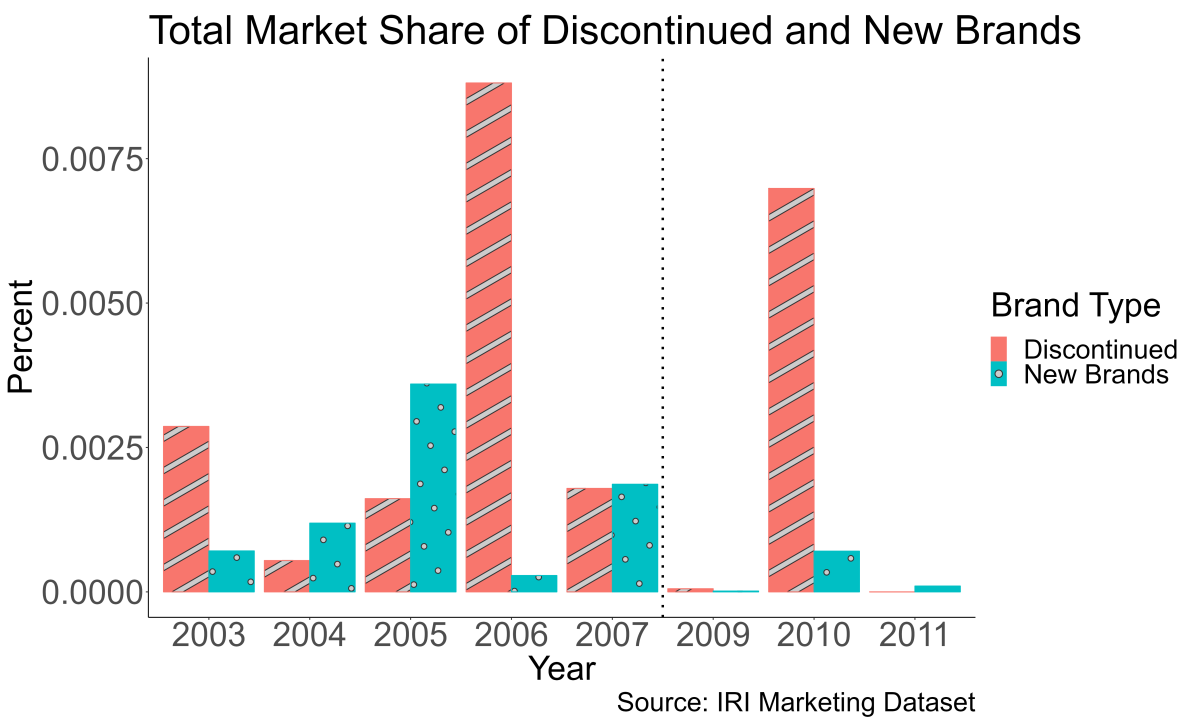

Note: This graph shows information from brands that were added in that year and brands that were discontinued after that year. 2008 is removed due to data issues. Prior to 2009, "new brands" and "discontinued brands" are the sum of the Coors and Miller new and discontinued brands, respectively. From 2009 onward, the "new brands" and "discontinued brands" are for MillerCoors only.

Figure 1.17: Mean Market Share of Discontinued and New Brands

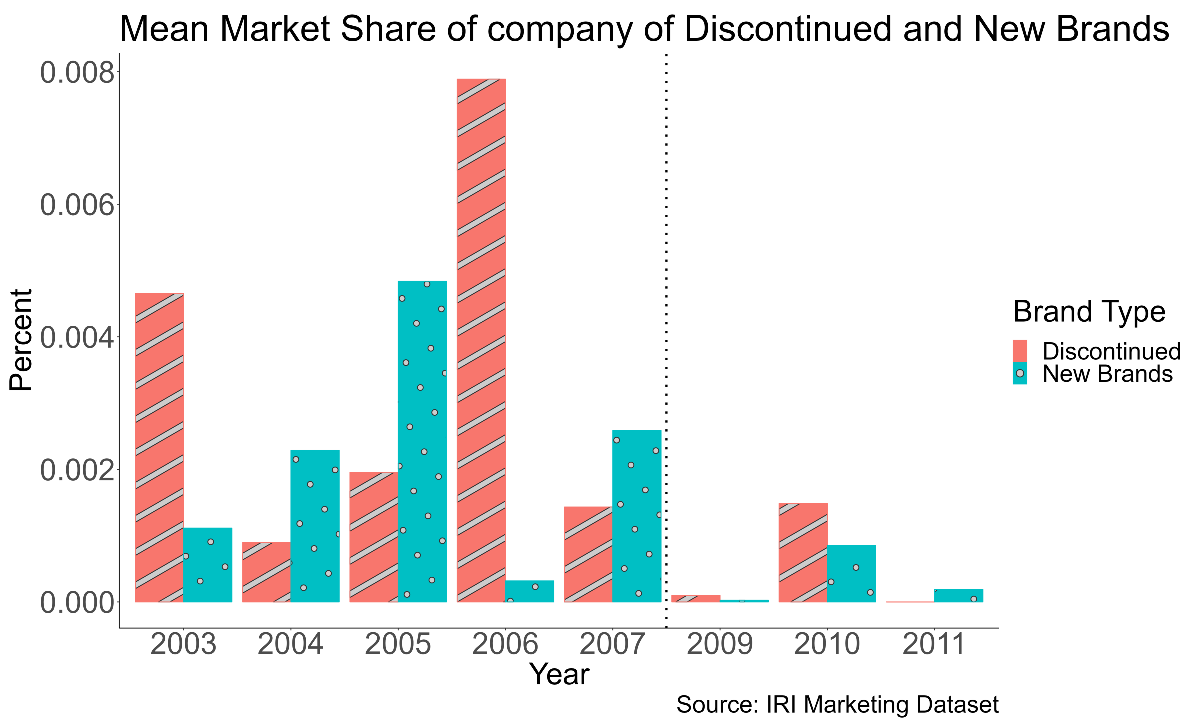

Note: This graph shows information from brands that were added in that year and brands that were discontinued after that year. 2008 is removed due to data issues. Prior to 2009, "new brands" and "discontinued brands" are the sum of the Coors and Miller new and discontinued brands, respectively. From 2009 onward, the "new brands" and "discontinued brands" are for MillerCoors only.

Figure 1.18: Number of Stores Carrying Discontinued and New Brands

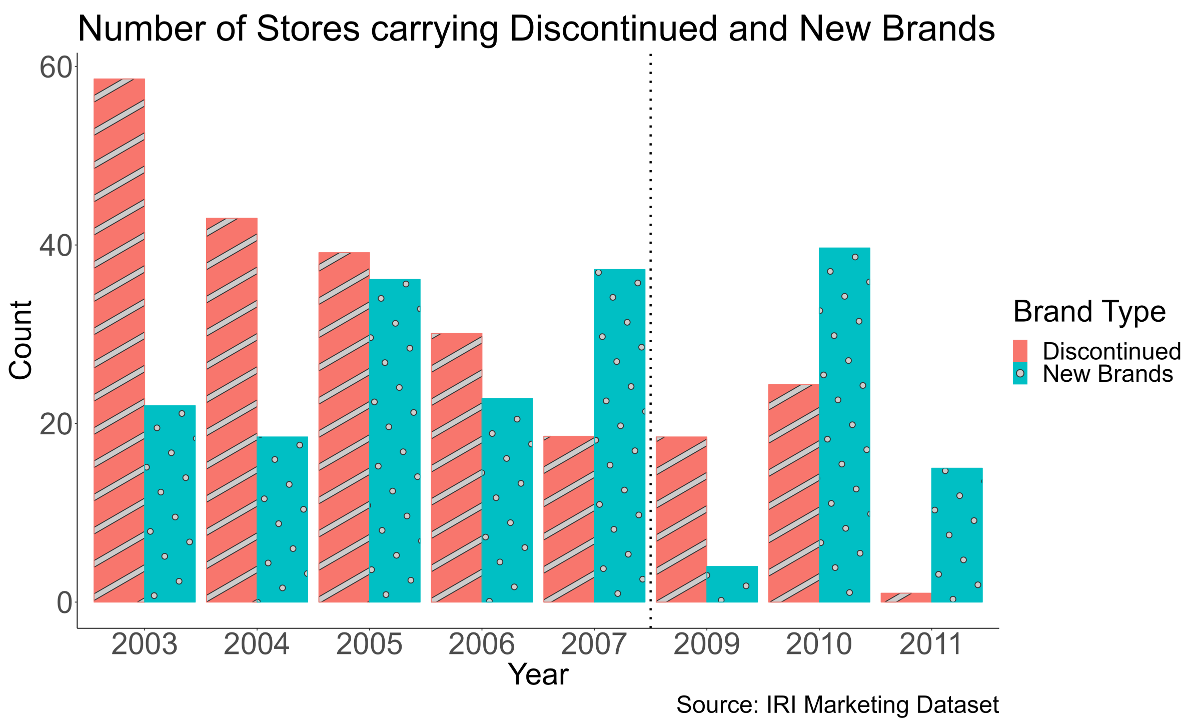

Note: This graph shows information from brands that were added in that year and brands that were discontinued after that year. 2008 is removed due to data issues. Prior to 2009, "new brands" and "discontinued brands" are the sum of the Coors and Miller new and discontinued brands, respectively. From 2009 onward, the "new brands" and "discontinued brands" are for MillerCoors only.

Figure 1.19: Number of Markets Where Discontinued and New Brands Were Sold

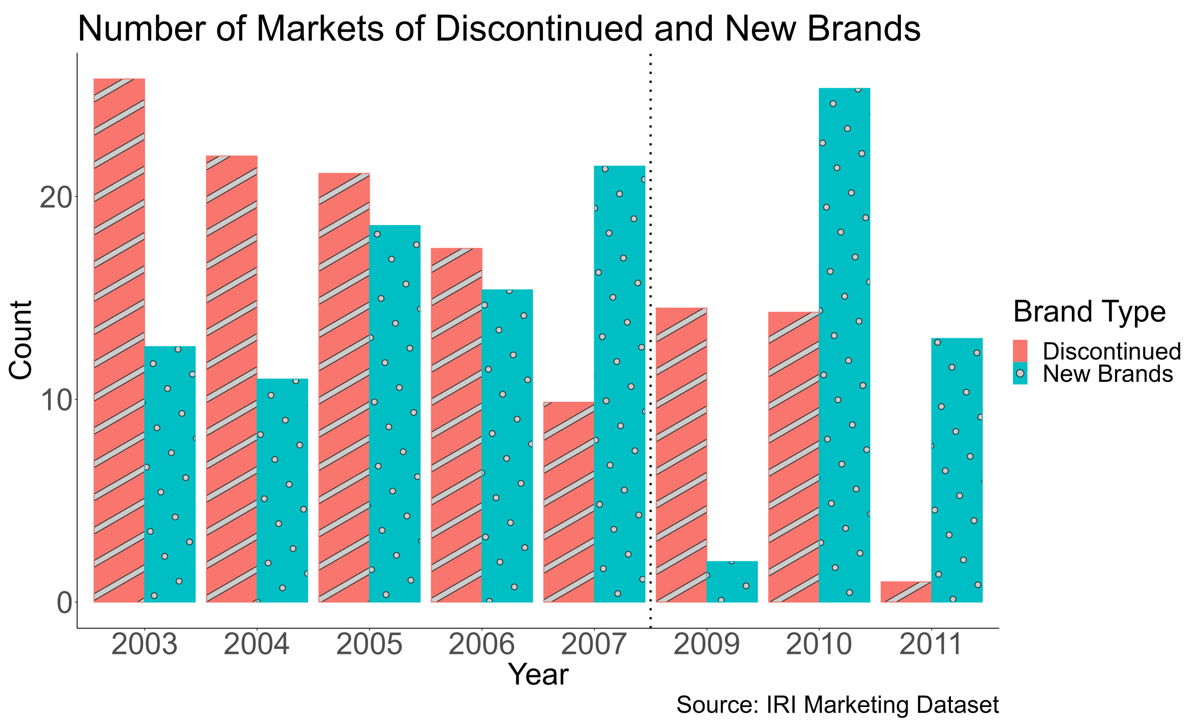

Note: This graph shows information from brands that were added in that year and brands that were discontinued after that year. 2008 is removed due to data issues. Prior to 2009, "new brands" and "discontinued brands" are the sum of the Coors and Miller new and discontinued brands, respectively. From 2009 onward, the "new brands" and "discontinued brands" are for MillerCoors only.

Overall, the data appears to be quite noisy prior to the merger but levels out for the most part after the merger. Notable trends from these figures are the changes in revenue, market share and sold products right before the merger, and the trend in market share after the merger. First, there were major declines prior to the merger in the number of sold products, but products introduced in that year made enough revenue to offset those losses, and similarly market share was high enough to offset the losses. However, after the merger, in 2010 and 2011, the discontinued products had more market share and more lost revenue than new products at the time. These trends are similar to the trends in new and discontinued brands in the paper with a slight decline after the merger leading to an increase in the last year.

Finally, these are not major products. This can be shown in the Mean Market Share Figure, which shows the share that new and discontinued products contribute to the firm's overall market share. At most, the new products contribute to 0.5% of their company's market share, and the discontinued products contribute to 0.3% of their company's market share.

The Beverage Marketing Corporation Brewery and Importer Database

The Beverage Marketing Corporation Brewery Database contains information on 554 select brewery Locations from 2006 to 2010. This dataset includes information on company names, location, key personnel, capacity, employment, location they serve, manufacturing lines, and what brands they produce at each brewery. This data can be linked to the the IRI dataset via Parent Company. Since multiple breweries make one brand, it is difficult to link the data to one brewery without making assumptions on what brewery makes what specific product shipped to the store.

Additionally, there is an importer supplement of the data that contains limited information on importers and imported brands. One key issue is that the importer data includes

Table 1.8: BMC Database Summary Statistics, at Individual Brewery Level

	Count			
Number of Unique Breweries	130			
Number of Unique Breweries, Coors	3			
Number of Unique Breweries, Miller	6			
Number of Firms in the top 5% percentile that sell imports	8			
Share of Firms in the top 5% percentile that sell imports and domestic brands	40%			
	Mean	Standard Deviation	Minimum	Maximum
Capacity, All Breweries	2425783.16	4230236.58	499.5	20000000
Capacity, Coors	11714285.43	7750576.28	5499999.5	20000000
Capacity, Miller	8166666.17	970142.50	6499999.5	9499999.5
Employment, All Breweries	327.86	611.13	2	3000
Employment, Coors	1567.57	1341.34	449.5	3000
Employment, Miller	791.17	121.27	624.5	874.5
Employment, All Importers	2425783.16	4230236.58	499.5	20000000
Employment, Coors Importers	3000.00	0	3000	3000
Employment, Miller Importers	11714285.43	7750576.28	5499999.5	20000000
# of Brands produced at each, All Breweries	9.68	5.65	1	20
# of Brands produced at each, Coors	8.86	6.4918	2	18
# of Brands produced at each, Miller	9.89	1.18	8	12
# of Brands produced at each, All Importers	1432.22	1395.84	112	3000
# of Brands produced at each, Coors Importers	3000.00	0.00	3000	3000
# of Brands produced at each, Miller Importers	8166666.17	970142.50	6499999.5	9500000

wine and liquor importers, with limited data on beer importers. Due to this constraint, I limit the importer data to only parent companies in the top 5 percentile of market share, as those can be clearly defined and located in the data. Other import beers are either not included in the IRI data or difficult to link due to brand name uncertainty. Importer data does not contain information on where capacity and the importer's home country, but does contain all other variables in the brewery database. I again chose the midpoint and topcodes for variables that only provided data in ranges.

For data cleaning, one important clarification is that the employment and capacity data is given in ranges and is topcoded. Due to this, I choose the midpoint of these ranges for summary statistics and estimation, and choose the value the data is topcoded at for topcodes. Table 1.8 shows the summary statistics for Capacity, Employment and Location for All breweries and Miller and Coors breweries and importers. Table 1.9 aggregates the breweries together at the parent company level to show summary statistics for total employment, capacity, brands and manufacturing lines.

Table 1.9: BMC Database Summary Statistics, at Parent Company Level

	Mean	Standard Deviation	Minimum	Maximum
All Breweries Number Breweries	1.26	1.43	1	13
All Breweries Total Brands	10.08	8.75	1	85
All Breweries Total Canlines	0.88	4.56	0	42
All Breweries Total Capacity	2232579.19	13042152.03	0	107999994
All Breweries Total Carbonated Bottling Lines	1.34	5.05	0	45
All Breweries Total Employment	376.18	1576.35	0	14269.5
Coors Number Breweries	2.33	0.58	2	3
Coors Total Brands	21.33	1.15	20	22
Coors Total Canlines	2.67	0.58	2	3
Coors Total Capacity	27333332.67	3175426.19	25499999.5	30999999
Coors Total Carbonated Bottling Lines	2.0000	0.0000	2	2
Coors Total Employment	3657.67	360.56	3449.5	4074
Miller Number Breweries	6.0000	0.0000	6	6
Miller Total Brands	21.67	2.31	19	23
Miller Total Canlines	16	0.0000	16	16
Miller Total Capacity	48999997.0000	0.0000	48999997	48999997
Miller Total Carbonated Bottling Lines	18.0000	0.0000	18	18
Miller Total Employment	4747.0000	0.0000	4747	4747

How Miller and Coors Breweries Compare to Other Breweries

Here I provide distributions of the summary statistic variables above and examine how Miller and Coors compare to other breweries in the dataset. These include employment, capital factors, and location of these breweries. Overall, I find that Miller and Coors were large breweries with large amounts of capital and labor, but lacked regional coverage. This provides evidence into the descriptive reasons for the merger as described in the paper.

Figure 1.20 shows the total employment for Miller and Coors before and after the merger, compared to its competitors. Overall, employment was on the higher end, yet slightly decreased due to the loss of the Tennessee Plant prior to the merger. There does not seem to be any other change in employment before and after the merger for Miller and Coors. Employment remained in the upper distribution of employment at most plants.

Figure 1.21 shows the capacity for Miller and Coors before and after the merger, compared to its competitors. Similar to employment, Miller and Coors are on the higher end of the distribution for capacity, with the majority of their breweries above six million barrels. Most

Figure 1.20: Brewery Employment, Miller and Coors

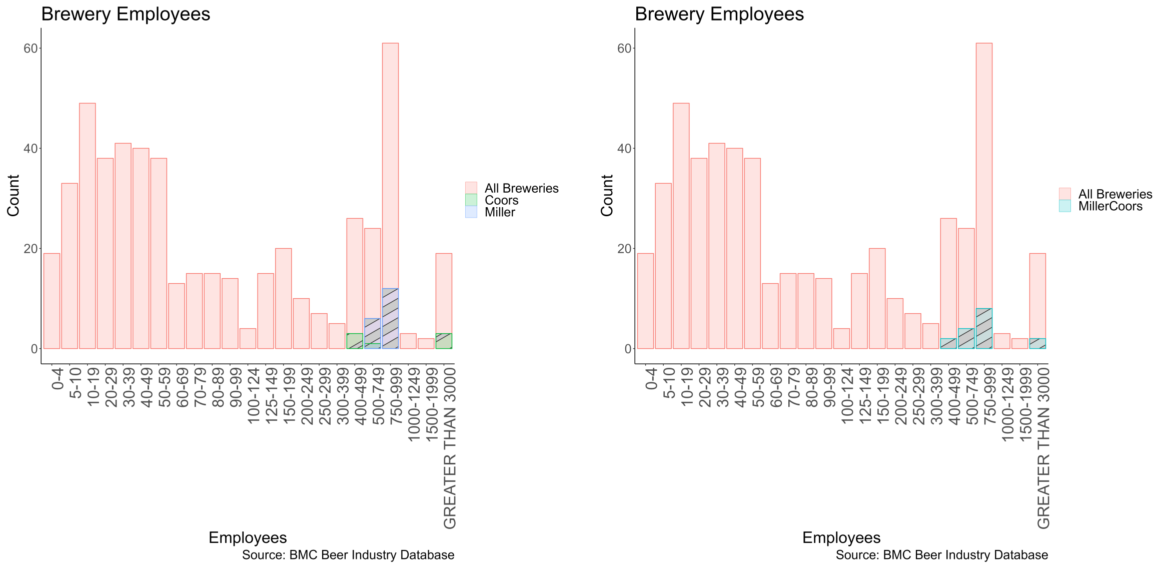

a) b)

Note: These figures show the total employment at each Miller and Coors brewery, compared to the total employment at each other brewery in the BMC dataset. The Memphis, TN Plant shut down in 2006, so there is one less MillerCoors plant post merger.

competitors fall far beneath that, with most capacities in the ten thousand to five hundred thousand range. A similar pattern holds for Figure 1.22, which shows that Miller and Coors were mostly in the upper distribution of canning and carbonated bottling lines among breweries.

Overall, these figures indicate Miller and Coors had a clear capital and employment advantage over most competitors. Regarding product variety, these companies had enough capacity and employees to sustain new varieties, and it also does not appear that the merger affected any of these aggregate variables that could have affected product variety.

Finally, Figure 1.23 compares the regional differences between Miller, Coors and their competitors. Two facts can be drawn from these figures. First, Coors and Miller had no overlap in region, which corroborates with comments made by regulators and spokespeople regarding reasons for the merger. Secondly, most breweries are located in the Northeast and the Pacific, and neither Miller or Coors have significant presence in the area. Only Miller

Figure 1.21: Brewery Capacity, Miller and Coors

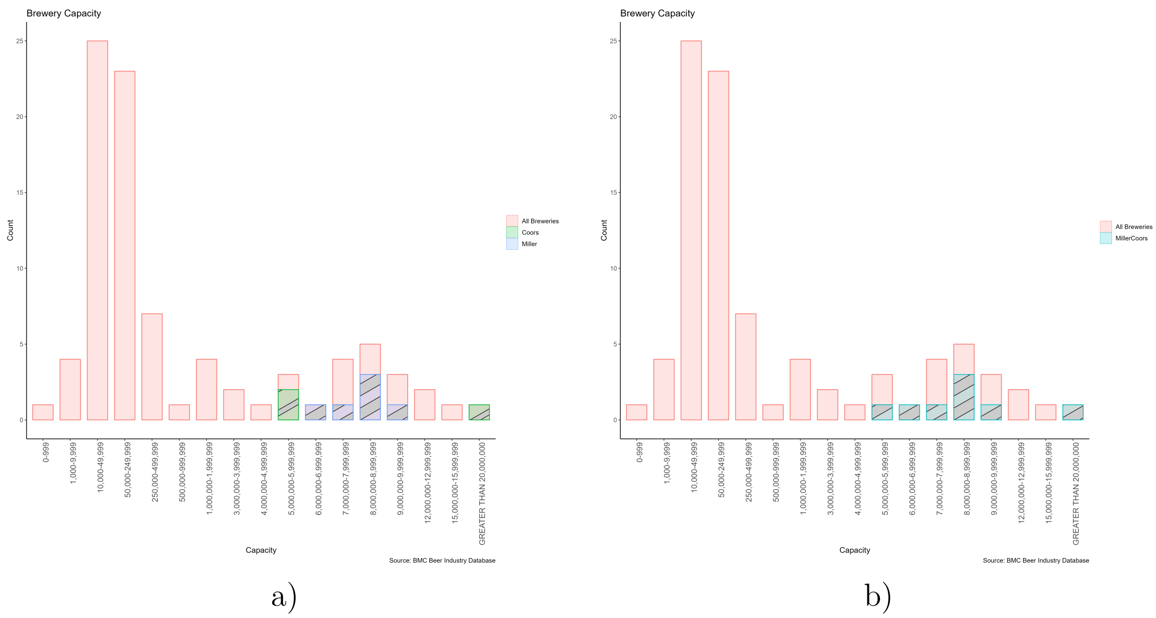

Note: These figures show the total capacity at each Miller and Coors brewery, compared to the total capacity at each other brewery in the BMC dataset. Capacity is measured in barrels, typically, 31 Gallons or approximately 117 liters. The Memphis, TN Plant shut down in 2006, so there is one less MillerCoors plant post merger.

has one brewery there.

In sum, is likely that the merger happened due to these regional differences, and had little effect on aggregate capital and labor trends at the breweries and elsewhere. Therefore, any decisions on product variety were not due to changes at the aggregate level of capital and labor, but due to changes in distance, changes in distribution costs or market power.

Differences Between the IRI dataset and the BMC Database

The IRI dataset and the BMC dataset differ in coverage due to the lack of categorization of parent companies in the BMC dataset. The BMC dataset's brewery database identifies breweries based on their name, while the IRI dataset identifies breweries based on their parent company and main vendor. As such, some brands that are subsidiaries of one company may not be listed as belonging to that company in the BMC database. Additionally, some

Figure 1.22: Comparison of Brewery Canlines Between Miller and Coors and All Other Firms in BMC Dataset

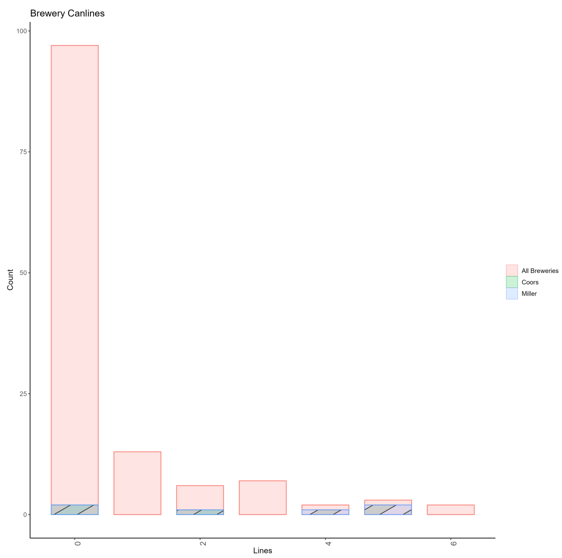

a)

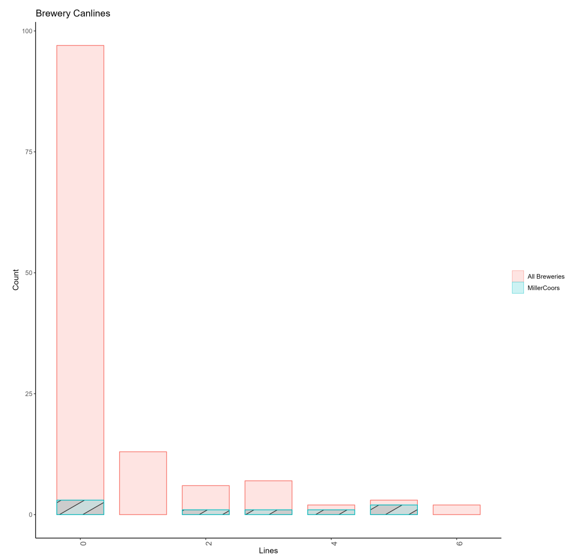

b)

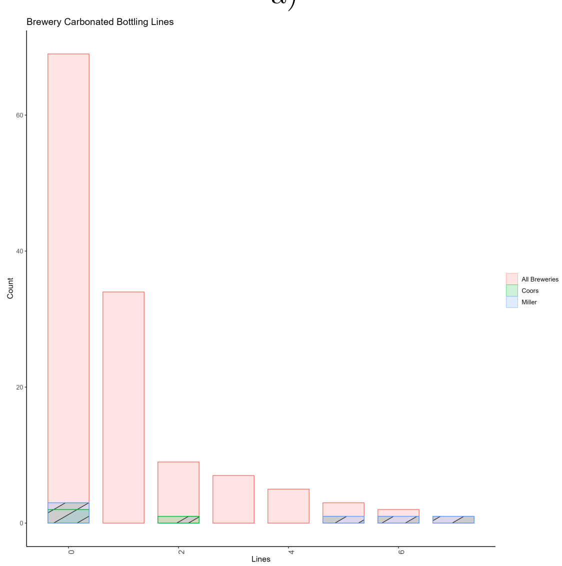

c)

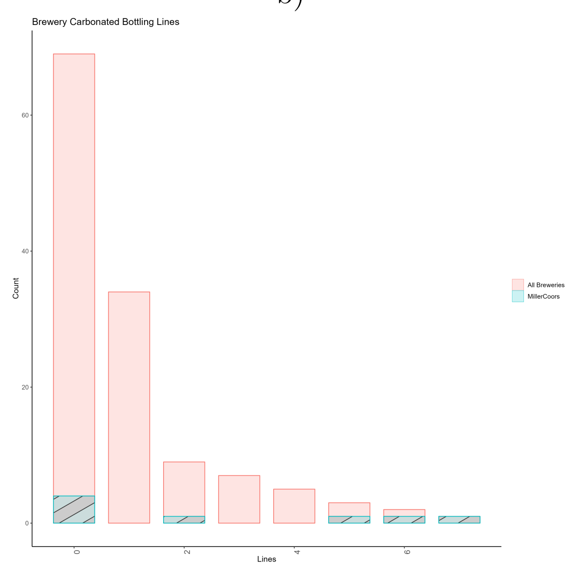

d)

Note: These figures show the total canlines and carbonated canlines at each Miller and Coors brewery, compared to the total canlines and carbonated canlines at each other brewery in the BMC dataset. The Memphis, TN Plant shut down in 2006, so there is one less MillerCoors plant post merger.

Figure 1.23: Comparison of Brewery Regions Between Miller and Coors and All Other Firms in BMC Dataset

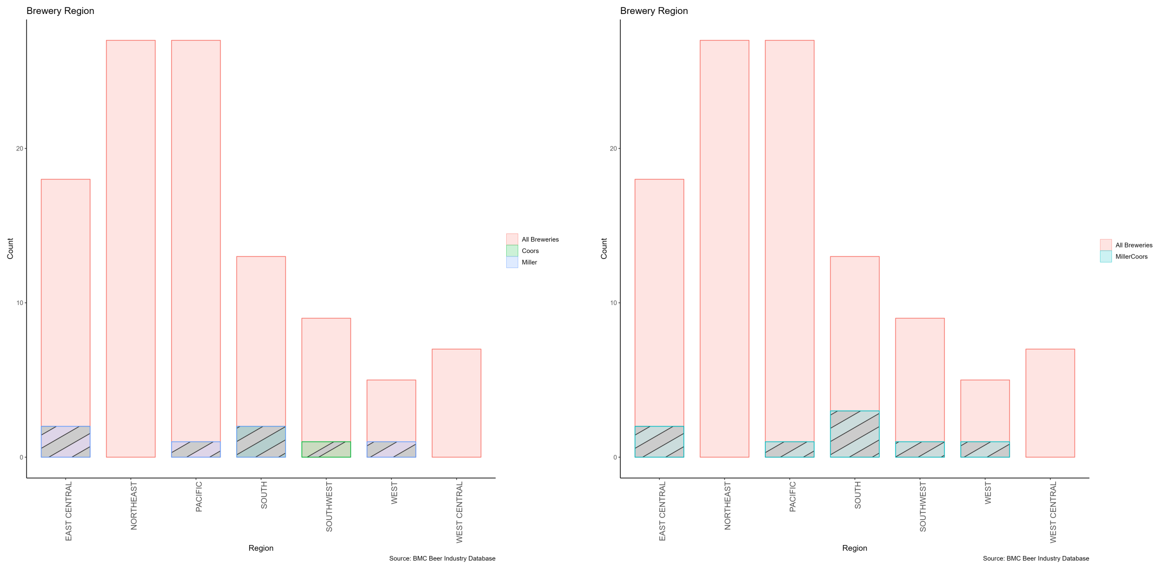

a) b)

Note: These figures show where each Miller and Coors brewery is located, compared to where each other brewery in the BMC dataset is located. The Memphis, TN Plant shut down in 2006, so there is one less MillerCoors plant post merger.

subsidiary breweries are not in the BMC database, making it impossible to verify whether these brands were produced in the current time period. This mainly affects Miller, as the majority of their brands are from subsidiary companies. Table 1.10 shows the difference between the number of brands in the BMC database versus the IRI database per year and the percentage of brands that were considered new or discontinued in the IRI dataset that were not in the BMC database.

The lack of coverage in the BMC database ranges from 25% to 100%, making the BMC database for brands difficult to use for this chapter. This is especially an issue for Miller, which have multiple subsidiaries not included in the BMC data, including Henry Weinhard, Mickey's, Hamm's, and Foster's. This change is further exacerbated by the merger, where the majority of products that leave the market are products from Miller subsidiaries. Due to this change, I use the IRI database for the majority of my analysis and use the brewery

Table 1.10: Difference Between Brand Coverage in the BMC Database and the IRI Database

Year	Company	Missing Discontinued Brands	Missing New Brands	Percent of Discontinued Brands in IRI	Percent of New Brands in IRI
2007	Coors	2	5	-25%	42%
2008	Coors	3		-33%	
2009	Coors	2		-25%	
2010	Coors	4		-40%	
2007	Miller	6	7	-75%	58%
2008	Miller	6		-67%	
2009	Miller	4		-50%	
2010	Miller	5	1	-50%	50%
2008	MillerCoors	9	10	-100%	100%
2009	MillerCoors	8	1	-100%	100%
2010	MillerCoors	10	2	-100%	100%

data to provide general information on production.

The Beverage Marketing Corporation Distributor Database

The Beverage Marketing Corporation Distributor data contains 9507 total distributors from 2006 to 2010. This dataset contains information on location, key personnel, parent companies they distribute from, employee count, sales count, and number of trucks. This data can be linked to the IRI dataset through the market level, as through conversations with the data providers and distributor employees, most distributors work within the market and do not travel across markets. Again, the data is given in ranges and topcoded, so the same data cleaning choices from the other Beverage Marketing Corporation datasets are used. Table 1.11 shows summary statistics for all distributors, distributors that distribute Coors products, and distributors that distribute Miller products.

1.6.2 Alternate Specifications and Robustness Tests of the Difference-in-Differences Model

In this section, I describe an alternate specification where I estimate the model on the raw count of product variety, rather than log transformations. I also describe the two robustness checks for the difference-in-differences results: A synthetic control for the changes in Miller,

Table 1.11: Summary Statistics, BMC Distributor Database

	Count			
Number of Unique Distributors	2080			
Number of Unique Distributors, Coors	826			
Number of Unique Distributors, Miller	769			
	Mean	Standard Deviation	Minimum	Maximum
All Distributors Employment	45.01	34.3216	1	100
All Distributors Sales	23781749.96	37129620.44	250000	200000000
All Distributors Trucks	17.88	19.41	1	100
Coors Distributors Employment	43.28	35.02	1	100
Coors Distributors Sales	22334800.86	36969137.02	250000	200000000
Coors Distributors Trucks	17.74	20.17	1	100
Miller Distributors Employment	42.38	35.02	1	100
Miller Distributors Sales	21626328.38	36119769.83	250000	200000000
Miller Distributors Trucks	17.33	19.98	1	100

Note: This table shows summary statistics for distributors from 2006 to 2010.

Coors and MillerCoors varieties, and a placebo test to verify whether the merger's effect is due to the merger itself and not a spurious pretrend. I find the model is robust to these changes and robustness tests. An additional test, that of using de-trended data that removes potential annual trends prior to the merger, is included in the paper.

Difference-in-Differences Results Without Logs

The following section describes a different version of the difference-in-differences model and the results. Similar to the paper's model, this uses a subset of the firm's data alongside the Miller and Coors data to estimate how MillerCoors brand offerings compare in regards to competitors. Unlike the model in the paper, the dependent variable is the raw number of brands, rather than the log transformation. I use the model below to estimate the impact of the merger on the number of brands of each firm i, in each market m at each period t:

$$\text{num brands}_{imt} = \beta_1(\text{Post Merger}_t) + \beta_2(\text{Miller}_i + \text{Coors}_i) + \beta_3(\text{MillerCoors}_i) + \epsilon_{imt} \quad (1.4)$$

Where (Post Merger_t) is an indicator for whether the observation is after the completion of the merger, $\text{Miller}_i + \text{Coors}_i$ is a sum of the indicators for Miller and Coors brand, and MillerCoors_i is an indicator for whether the brand is a MillerCoors brand. The last coefficient acts as a difference-in-differences coefficient of interest, as it measures the additional impact of the number of brands after the merger and under the merged company, compared to the control group. As before, I employ two separate control groups and estimate the model for each.

Overall, the difference-in-differences coefficient is positive for the first group but null for the second group. Table 1.12 shows the difference-in-differences result when comparing against all products in the market. Here, the coefficient is 3.56 under all controls, implying that after the merger, MillerCoors gained 3.56 brands relative to all products. Comparisons between all breweries in the market and MillerCoors have potential issues relating to unobsevables, as described in the paper. Table 1.13 shows that for the second group, the difference-in-differences coefficient is indistinguishable from zero for every specifications of the model except the first. Given that the model in the paper ranges from a positive coefficient to a negative one under year fixed effects, this points towards a potential annual trend driving these results. I correct for this in the detrended models in the paper.

Synthetic Control

The following section creates a synthetic control group for MillerCoors using the methods of Abadie and Gardeazabal (2003) to better compare product variety changes of the firm after the merger. This method uses a combination of control group units to create a better comparison to the treated group. In the context of this work, I create a convex combination of the breweries in the top five percentile to create a better control group for Miller, Coors and MillerCoors.

One issue with this method is that the control group differs greatly on levels. As shown

Table 1.12: Difference-in-Differences Results, Number of Brands Only

	Dependent variable:				
	Number of Brands				
	(1)	(2)	(3)	(4)	(5)
Post Merger	5.006^{***}	0.585^{***}	0.545^{***}	-0.142^{***}	-0.164^{***}
	(0.093)	(0.083)	(0.078)	(0.046)	(0.049)
Miller or Coors	38.835^{***}	34.410^{***}	34.395^{***}	34.413^{***}	34.414^{***}
	(0.904)	(0.884)	(0.884)	(0.885)	(0.885)
MillerCoors	-0.880	3.544^{***}	3.585^{***}	3.564^{***}	3.564^{***}
	(0.613)	(0.613)	(0.607)	(0.604)	(0.604)
State FE	No	No	Yes	Yes	Yes
Month FE	No	Yes	Yes	No	Yes
Year FE	No	No	No	Yes	Yes
Observations	298,153	298,153	298,153	298,153	298,153
Adjusted R^2	0.400	0.494	0.495	0.495	0.495
Note:			*p<0.1; **p<0.05; ***p<0.01		

Note: The table above shows the difference-and-difference results under a non-logged dependant variable. All other estimation is conducted the same. Observations are at the firm–market–month–year level. The “Post Merger” variable indicates the period after July 2008, the start of the merged company’s operation. All firms that product beer products in the IRI dataset are included in this regression. Standard errors clustered at the market level. For the last set of regressions, to remove colinearity with the year, month and Post Merger variable, the December fixed effect is removed.

Table 1.13: Difference-in-Differences Results for Top 5% of Market Share Firms, Number of Brands Only

	Dependent variable:				
	Number of Brands				
	(1)	(2)	(3)	(4)	(5)
Post Merger	16.519***	3.911***	3.901***	0.211***	−0.074
	(0.306)	(0.217)	(0.214)	(0.081)	(0.091)
Miller or Coors	38.835***	26.217***	26.214***	26.244***	26.242***
	(0.904)	(0.870)	(0.862)	(0.861)	(0.861)
MillerCoors	−12.393***	0.225	0.235	0.210	0.213
	(0.476)	(0.587)	(0.585)	(0.585)	(0.585)
State FE	No	No	Yes	Yes	Yes
Month FE	No	Yes	Yes	No	Yes
Year FE	No	No	No	Yes	Yes
Observations	45,712	45,712	45,712	45,712	45,712
Adjusted R^2	0.524	0.667	0.672	0.674	0.674
Note:			*p<0.1; **p<0.05; ***p<0.01		

Note: The table above shows the difference-and-difference results under a non-logged dependant variable. All other estimation is conducted the same. Observations are at the firm–market–month–year level. The "Post Merger" variable indicates the period after July 2008, the start of the merged company's operation. Only the ten largest firms by national revenue share prior to 2007 are included in this regression. Standard errors clustered at the market level. For the last set of regressions, to remove colinearity with the year, month and Post Merger variable, the December fixed effect is removed.

in the figure in the paper, Anheuser-Busch Inbev is the most appropriate comparison group prior to the merger, so any synthetic control should use Anheuser-Busch Inbev. One solution could be to split MillerCoors into separate companies, so the synthetic control can better match both Miller and Coors on levels. This causes additional issues, as the following figures representing the synthetic control on just Miller, just Coors and MillerCoors show.

Figure 1.24 shows the synthetic control for Miller and Miller's observed outcomes. The post-period Miller observations are extrapolated by calculating the average share of Miller goods contributing to the simulated MillerCoors number of brands prior to the merger. Overall, the Miller synthetic control matches the pre-period fairly well, and the figure does suggest a positive relationship of product variety after the merger. However, Figure 1.25 and Figure 1.26 shows the synthetic control comparison for Coors and MillerCoors, respectively. Although these match better on levels than the control group, the large fluctuations prior to the merger are concerning for the parallel trends assumption.

In all three figures, the synthetic control is consistently above the number of brands observed by Miller, Coors and MillerCoors. With the synthetic control, I still observe the number of brands is less than what would be expected, so the merger still has a negative effect on the number of brands. As an additional statistic, I calculate the average difference after the merger between the observed and synthetic brands for all three scenarios. I find that the average difference for Miller and its synthetic counterpart, Coors and its synthetic counterpart and MillerCoors and its synthetic counterpart is 4.82 brands, 2.98 brands and 7 brands, respectively. This effect is lower than the difference-in-difference estimates of about 10 brands, but still shows a decline in the number of brands relative to their competitors.

Placebo Test

Here, I describe the Placebo Test, where I examine whether the same results are observed if the merger was moved July 2003 rather than July 2008. This is to address concerns that the

Figure 1.24: Number of Brands for Miller, Synthetic and Actual

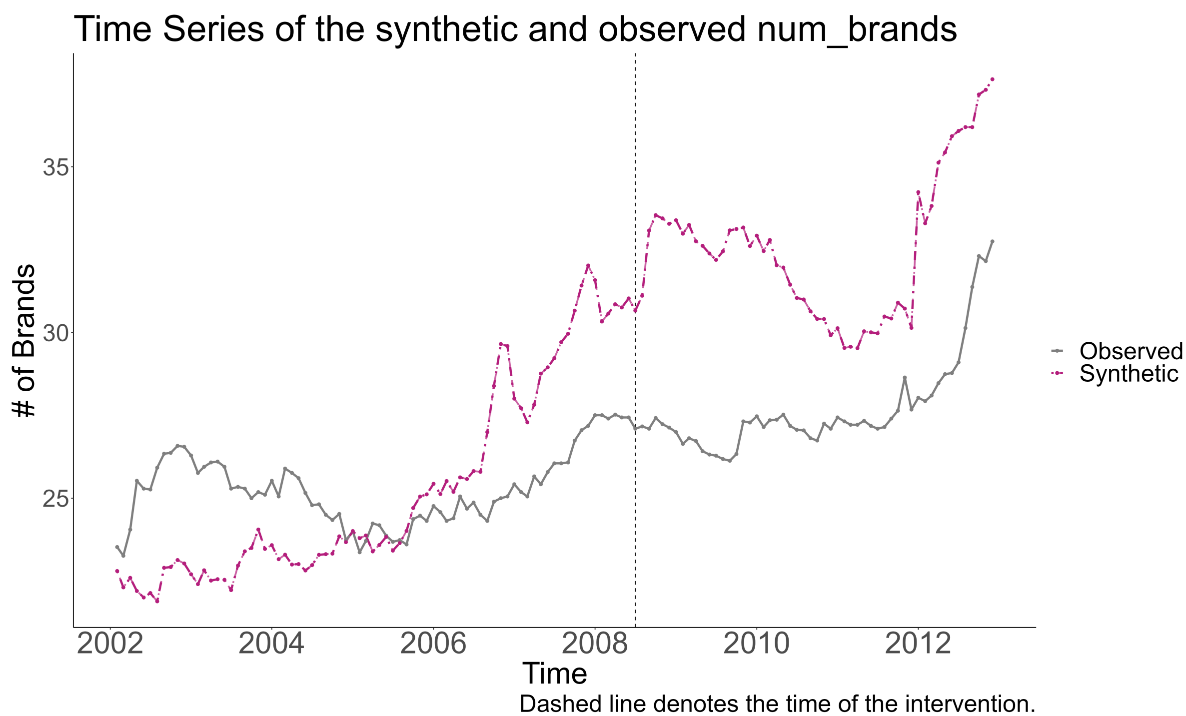

Note: This graph plots the average of the number of brands offered over all markets for Miller, and the synthetic control for Miller using firms within the top 5 percentile of market share. The Red dashed line denotes the date the MillerCoors merger was finalized. Miller is extrapolated to future periods by calculating the share that Miller goods contributed to MillerCoors prior to the merger, then taking the average. This value is approximately 0.63. Weights for the synthetic control are as follows: 39% for Anheuser-Busch, 60% for S&P Company, and negligible amounts for Boston Beer Company, Constellation Brands Inc, Yuengling & Son Inc, Deschutes Brewery, Diageo, and Heineken USA.

Figure 1.25: Number of Brands for Coors, Synthetic and Actual

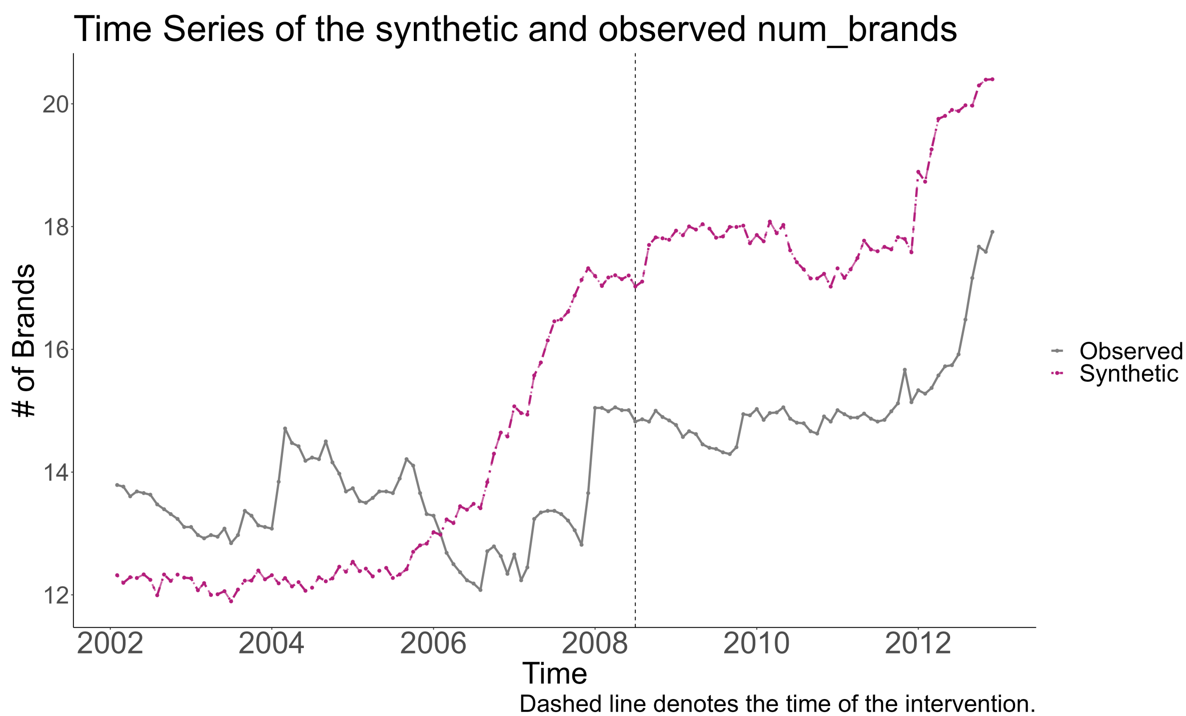

Note: This graph plots the average of the number of brands offered over all markets for Coors, and the synthetic control for Miller using firms within the top 5 percentile of market share. The Red dashed line denotes the date the MillerCoors merger was finalized. Coors is extrapolated to future periods by calculating the share that Miller goods contributed to MillerCoors prior to the merger, then taking the average. This value is approximately 0.36. Weights for the synthetic control are as follows: 70% for S&P Company, 20% for Diageo, 7.6 % for Anheuser-Busch, and negligible amounts for Boston Beer Company, Constellation Brands Inc, Yuengling & Son Inc, Deschutes Brewery, and Heineken USA.

Figure 1.26: Number of Brands for MillerCoors, Synthetic and Actual

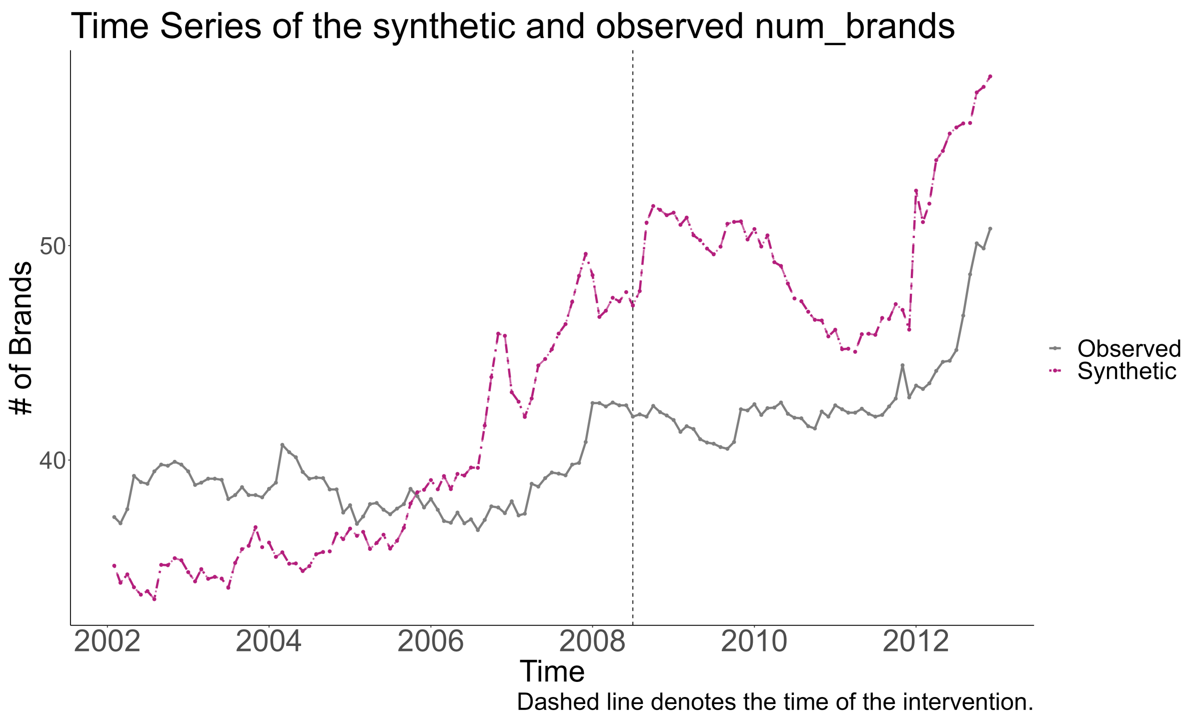

Note: This graph plots the average of the number of brands offered over all markets for MillerCoors, and the synthetic control for MillerCoors using firms within the top 5 percentile of market share. The Red dashed line denotes the date the MillerCoors merger was finalized. MillerCoors is extrapolated to past periods by summing Miller and Coors before the merger. Weights for the synthetic control are as follows: 86% for AB-Inbev and 2% for S&P Company, Diageo, Boston Beer Company, Constellation Brands Inc, Yuengling & Son Inc, Deschutes Brewery, and Heineken USA.

results I am seeing are coming from the merger itself, rather than other trends. To do this, I remove the post July 2008 data, then re-estimate the model having adjusted the coefficients such that the merger occurs in July 2003 instead. This means the post merger dummy in the model is coded to at and after July 2003, and the MillerCoors dummy is also coded to at and after this date. If the coefficient is significantly different than the coefficient observed in the main results, then the main results are not caused by a spurious pre-trend and provides evidence that the main results are due to the merger itself.

I find the two difference-in-differences models for both brands and product variety pass the placebo test and estimate coefficients that are outside of the bounds of the baseline results. While some of these have the same sign as the baseline results, I interpret this as the merger had a stronger impact on an existing negative pre-trend. Table 1.14 shows the results for the placebo test for MillerCoors compared to all firms on brands, and Table 1.15 shows the results for the placebo test for MillerCoors compared to all firms at the brand level and at the product level. Table 1.16 and Table 1.17 shows the results for the placebo test for MillerCoors compared to all top firms at the brand level and at the product level.

Table 1.14: All Breweries Difference-in-Differences, Placebo Test Results with Product Variety as Number of Brands

	Dependent variable:				
	log(Number of Brands)				
	(1)	(2)	(3)	(4)	(5)
Post Merger	1.005***	0.024*	0.023*	−0.008	−0.017**
	(0.016)	(0.014)	(0.014)	(0.009)	(0.008)
Miller or Coors	3.653***	2.670***	2.669***	2.669***	2.669***
	(0.022)	(0.023)	(0.023)	(0.023)	(0.023)
MillerCoors	−1.014***	−0.031**	−0.030**	−0.027**	−0.027**
	(0.018)	(0.012)	(0.012)	(0.012)	(0.012)
State FE	No	No	Yes	Yes	Yes
Month FE	No	Yes	Yes	No	Yes
Year FE	No	No	No	Yes	Yes
Observations	146,238	146,238	146,238	146,238	146,238
Adjusted R^2	0.516	0.614	0.615	0.616	0.616
Note:					*p<0.1; **p<0.05; ***p<0.01

Note: The table above shows the results of the placebo test for the MillerCoors merger at the brand level. To conduct the placebo test, I remove all data at and after the merger's finalization date of July 2008, then set the Post Merger Dummy and the MillerCoors dummy to July 2003. These results show the comparison to all firms in the dataset as described by the paper. Observations are at the firm–market–month–year level. Standard errors clustered at the market level. For the last set of regressions, to remove colinearity with the year, month and Post Merger variable, the December fixed effect is removed.

Table 1.15: All Breweries Difference-in-Differences, Placebo Test Results with Product Variety as Number of Products

	Dependent variable:				
	log(Number of Products)				
	(1)	(2)	(3)	(4)	(5)
Post Merger	1.305^{***}	0.014	0.013	0.005	-0.001
	(0.019)	(0.017)	(0.016)	(0.010)	(0.010)
Miller or Coors	4.836^{***}	3.545^{***}	3.545^{***}	3.545^{***}	3.545^{***}
	(0.029)	(0.032)	(0.033)	(0.033)	(0.033)
MillerCoors	-1.326^{***}	-0.035^{*}	-0.034^{*}	-0.033^{*}	-0.033^{*}
	(0.024)	(0.018)	(0.018)	(0.018)	(0.017)
State FE	No	No	Yes	Yes	Yes
Month FE	No	Yes	Yes	No	Yes
Year FE	No	No	No	Yes	Yes
Observations	146,238	146,238	146,238	146,238	146,238
Adjusted R^2	0.540	0.644	0.646	0.646	0.646
Note:			*p<0.1; **p<0.05; ***p<0.01		

Note: The table above shows the results of the placebo test for the MillerCoors merger at the product level. To conduct the placebo test, I remove all data at and after the merger's finalization date of July 2008, then set the Post Merger Dummy and the MillerCoors dummy to July 2003. These results show the comparison to all firms in the dataset as described by the paper. Observations are at the firm–market–month–year level. Standard errors clustered at the market level. For the last set of regressions, to remove colinearity with the year, month and Post Merger variable, the December fixed effect is removed.

Table 1.16: Top 5% Market Share Difference-in-Differences, Placebo Test Results with Product Variety as Number of Brands

	Dependent variable:				
	Number of Brands				
	(1)	(2)	(3)	(4)	(5)
Post Merger	2.186***	0.086***	0.087***	−0.0003	−0.020**
	(0.024)	(0.017)	(0.017)	(0.007)	(0.008)
Miller or Coors	3.653***	1.549***	1.548***	1.548***	1.547***
	(0.022)	(0.031)	(0.031)	(0.031)	(0.030)
MillerCoors	−2.194***	−0.090***	−0.090***	−0.086***	−0.086***
	(0.021)	(0.014)	(0.014)	(0.014)	(0.014)
State FE	No	No	Yes	Yes	Yes
Month FE	No	Yes	Yes	No	Yes
Year FE	No	No	No	Yes	Yes
Observations	25,577	25,577	25,577	25,577	25,577
Adjusted R^2	0.765	0.905	0.906	0.907	0.907
Note:					*p<0.1; **p<0.05; ***p<0.01

Note: The table above shows the results of the placebo test for the MillerCoors merger at the brand level. To conduct the placebo test, I remove all data at and after the merger's finalization date of July 2008, then set the Post Merger Dummy and the MillerCoors dummy to July 2003. These results show the comparison to the ten largest firms by national revenue share prior to 2007 as described by the paper. Observations are at the firm–market–month–year level. Standard errors clustered at the market level. For the last set of regressions, to remove colinearity with the year, month and Post Merger variable, the December fixed effect is removed.

Table 1.17: Top 5% Market Share Difference-in-Differences, Placebo Test Results with Product Variety as Number of Products

	Dependent variable:				
	Number of Products				
	(1)	(2)	(3)	(4)	(5)
Post Merger	2.928***	0.136***	0.137***	0.020**	−0.007
	(0.029)	(0.021)	(0.021)	(0.009)	(0.011)
Miller or Coors	4.836***	2.040***	2.042***	2.042***	2.041***
	(0.029)	(0.034)	(0.034)	(0.034)	(0.034)
MillerCoors	−2.949***	−0.153***	−0.154***	−0.151***	−0.150***
	(0.027)	(0.020)	(0.020)	(0.020)	(0.020)
State FE	No	No	Yes	Yes	Yes
Month FE	No	Yes	Yes	No	Yes
Year FE	No	No	No	Yes	Yes
Observations	25,577	25,577	25,577	25,577	25,577
Adjusted R^2	0.776	0.916	0.918	0.918	0.918
Note:	*p<0.1; **p<0.05; ***p<0.01				

Note: The table above shows the results of the placebo test for the MillerCoors merger. To conduct the placebo test, I remove all data at and after the merger's finalization date of July 2008, then set the Post Merger Dummy and the MillerCoors dummy to July 2003. These results show the comparison to the ten largest firms by national revenue share prior to 2007 as described by the paper. Observations are at the firm–market–month–year level. Standard errors clustered at the market level. For the last set of regressions, to remove colinearity with the year, month and Post Merger variable, the December fixed effect is removed.

1.7 References

Abadie, Alberto, and Javier Gardeazabal. 2003. "The Economic Costs of Conflict: A Case Study of the Basque Country." *American Economic Review* 93 (March): 113–32. https://doi.org/10.1257/000282803321455188.

Ashenfelter, Orley C., Daniel S. Hosken, and Matthew C. Weinberg. 2015. "Efficiencies Brewed: Pricing and Consolidation in the US Beer Industry." *The RAND Journal of Economics* 46 (June): 328–61. https://doi.org/10.1111/1756-2171.12092.

Ashenfelter, Orley, Daniel Hosken, and Matthew Weinberg. 2014. "Did Robert Bork Understate the Competitive Impact of Mergers? Evidence from Consummated Mergers." *Https://Doi.org/10.1086/675862* 57 (August): S67–100. https://doi.org/10.1086/675862.

Asker, John. 2016 "DIAGNOSING FORECLOSURE DUE TO EXCLUSIVE DEALING*." https://doi.org/10.1111/joie.12114.

Association, Brewers. 2022. "National Beer Sales & Production Data." https://www.brewersassociation.org/statistics-and-data/national-beer-stats/.

Atalay, Enghin, Alan T. Sorensen, Christopher J. Sullivan, and Wanjia Zhu. 2023. "Product Repositioning by Merging Firms," May. https://doi.org/10.3386/W31229.

Barajas, Jesus M, Geoff Boeing, and Julie Wartell. 2017. *Untapped: Exploring the Cultural Dimensions of Craft Beer*.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica* 63 (July): 841. https://doi.org/10.2307/2171802.

Capobianco, Antonio. 2018. "Non-Price Effects of Mergers." *129th OECD Competition Committee Meeting*. https://www.oecd.org/competition/non-price-effects-of-mergers.htm.

Chandra, Ambarish, and Matthew Weinberg. 2018. “How Does Advertising Depend on Competition? Evidence from u.s. Brewing.” *Https://Doi.org/10.1287/Mnsc.2017.2889* 64 (January): 5132–48. https://doi.org/10.1287/MNSC.2017.2889.

Ciliberto, Federico, and Jonathan W Williams. 2014. “Does Multimarket Contact Facilitate Tacit Collusion? Inference on Conduct Parameters in the Airline Industry.” *Journal of Economics* 45: 764–91. https://doi.org/10.1111/1756-2171.12070.

Conlon, Christopher, and Jeff Gortmaker. 2020. “Best Practices for Differentiated Products Demand Estimation with PyBLP.” *RAND Journal of Economics* 51 (December): 1108–61. https://doi.org/10.1111/1756-2171.12352.

Conlon, Christopher T., and Nirupama L. Rao. 2020. “Discrete Prices and the Incidence and Efficiency of Excise Taxes."” American Economic Journal: Economic Policy 12 4 111-43 2020 10.1257/pol.20160391 https://www.aeaweb.org/articles?id=10.1257/pol.20160391

Dale, Stacy Berg, and Alan B. Krueger. 2002. “Estimating the Payoff to Attending a More Selective College: An Application of Selection on Observables and Unobservables.” *The Quarterly Journal of Economics* 117 (November): 1491–1527. https://doi.org/10.1162/003355302320935089.

Demirer, Mert, and Omer Karaduman. 2023. “Do Mergers and Acquisitions Improve Efficiency: Evidence from Power Plants,” January. https://gsb-faculty.stanford.edu/omer-karaduman/files/2022/12/Draft.pdf.

Division, Antitrust. 2008. “STATEMENT OF THE DEPARTMENT OF JUSTICE’s ANTITRUST DIVISION ON ITS DECISION TO CLOSE ITS INVESTIGATION OF THE JOINT VENTURE BETWEEN SABMILLER PLC AND MOLSON COORS BREWING COMPANY.”

Elzinga, Kenneth G., Carol Horton Tremblay, and Victor J. Tremblay. 2015. “Craft Beer in the United States: History, Numbers, and Geography*." *Journal of Wine Economics* 10 (December): 242–74. https://doi.org/10.1017/JWE.2015.22.

Fan, Ying, and Chenyu Yang. 2022. "Estimating Discrete Games with Many Firms and Many Decisions: An Application to Merger and Product Variety." *SSRN Electronic Journal*, May. https://doi.org/10.2139/SSRN.3681556.

Goldberg, Pinelopi Koujianou, and Rebecca Hellerstein. 2013. "A Structural Approach to Identifying the Sources of Local Currency Price Stability." *The Review of Economic Studies* 80 (January): 175–210. https://doi.org/10.1093/RESTUD/RDS015.

Goodman-Bacon, Andrew. 2021. "Difference-in-Differences with Variation in Treatment Timing." *Journal of Econometrics* 225 (December): 254–77. https://doi.org/10.1016/J.JECONOM.2021.03.014.

Grennan, Matthew. 2013. "Price Discrimination and Bargaining: Empirical Evidence from Medical Devices." *American Economic Review* 103 (February): 145–77. https://doi.org/10.1257/AER.103.1.145.

Grigolon, Laura, and Frank Verboven. 2014. "Nested Logit or Random Coefficients Logit? A Comparison of Alternative Discrete Choice Models of Product Differentiation." *The Review of Economics and Statistics* 96 (December): 916–35. https://doi.org/10.1162/REST_A_00420.

Hellerstein, Rebecca. 2008. "Who Bears the Cost of a Change in the Exchange Rate? Pass-Through Accounting for the Case of Beer." *Journal of International Economics* 76 (September): 14–32. https://doi.org/10.1016/J.JINTECO.2008.03.007.

Heyer, Ken, Carl Shapiro, Jeffrey Wilder, Ken Heyer, Carl Shapiro, and Jeffrey Wilder. 2009. "The Year in Review: Economics at the Antitrust Division, 2008-2009." *Review of Industrial Organization* 35 (December): 349–67. https://doi.org/10.1007/S11151-009-9232-1.

Khmelnitskaya, Ekaterina, Guillermo Marshall, and Scott Orr, 2024. "Identifying Scale and Scope Economies using Demand-Side Data". *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4757085

Martin, Andrew. 2007. "Merger for SABMiller and Molson Coors." *New York Times*. https://www.nytimes.com/2007/10/10/business/worldbusiness/10beer.html.

Mcfadden, Daniel. 1977. "Modelling the Choice of Residential Location." https://elischolar.library.yale.edu/cgi/viewcontent.cgi?article=1709&context=cowles-discussion-paper-series.

Miller, Nathan H., and Matthew C. Weinberg. 2017. "Understanding the Price Effects of the MillerCoors Joint Venture." *Econometrica* 85 (November): 1763–91. https://doi.org/10.3982/ECTA13333.

Petrin, Amil. 2002. "Quantifying the Benefits of New Products: The Case of the Minivan." *Journal of Political Economy*. Vol. 110.

Tremblay, Victor J., Carol Horton Tremblay, Victor J. Tremblay, and Carol Horton Tremblay. 2005. "The US Brewing Industry: Data and Economic Analysis" 1: 379. https://EconPapers.repec.org/RePEc:mtp:titles:0262201518.

Weinberg, Matthew, Gloria Sheu, and Nathan Miller. 2019. "Oligopolistic Price Leadership and Mergers: An Empirical Model of the u.s. Beer Industry." *2019 Meeting Papers*. https://ideas.repec.org/p/red/sed019/1210.html.

Wollmann, Thomas G. 2018. "Trucks Without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles." *American Economic Review* 108 (June): 1364–1406. https://doi.org/10.1257/AER.20160863.

Chapter 2

How Much Do Consumers Care About New and Discontinued Products? A Case Study of the MillerCoors Merger

2.1 Introduction

One issue that regulators face in merger analysis is how to quantify the nonprice effects of a merger and what their importance is relative to price effects. According to a recent OECD report summarizing the policies of twenty-one competition agencies around the world, eighteen explicitly have policies addressing nonprice effects of mergers, but the majority do not address these issues unless there are "claims made by merging parties, their customers/consumers, and rivals".[1] However, it is unclear how one such nonprice effect, that on product variety, should be weighted in the social optimum, as some agencies consider them second order to

[1] In such cases, the complaints may be addressed through qualitative evidence in the absence of direct measurement (Capobianco 2018).

price effects while others consider them of equal importance to price effects. Resolving this issue requires quantifying both the effect of the merger on product variety and the consumer welfare effect of product variety changes for direct comparison.

Several papers in the literature have provided models to quantify the effects of changing product variety but make assumptions on which products leave and enter markets. Wollmann (2018) studies welfare effects from product variety changes in a theoretical setting where a bailout never occurs and bankrupt firms are acquired by competitors. Fan and Yang (2022) examine how to model changes in product variety in multiple markets in a hypothetical where a major brewery acquires smaller craft breweries. Overall, these papers find a loss in consumer welfare from the decline in product variety. However, without an appropriate benchmark, it is unclear whether this effect should be a priority for policymakers relative to price effects from a merger.

In this chapter, I use the setting of the brewery industry after the MillerCoors merger of 2008 to estimate the consumer welfare effects of these changes, and describe how they compare to the consumer welfare effects of price changes measured in the literature. Prior work has not considered product variety changes and their potential effects on consumers in the context of this widely studied merger. In the first chapter, I describe the setting, show how price is limited in this market and discuss how much the product variety effects can be attributed to the merger itself. In this chapter, I estimate a model of consumer demand to measure the consumer welfare effects of the change in product variety. I find that in terms of consumer welfare, the gain from new product variety far outweighed the loss from discontinued product variety. For comparison, I find that the gain in consumer surplus from new products and the decline from discontinued products are approximately 34% and -4% of the consumer welfare effects, respectively, of coordinated pricing found in Miller and Weinberg (2017).

The intuition behind why consumer surplus changes matter can be motivated through

Figure 2.1: Visual Representation of Product Introduction

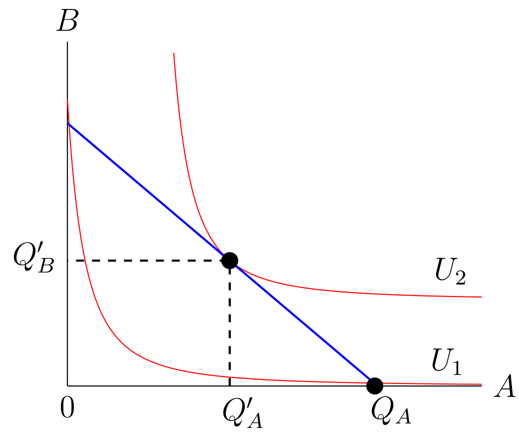

a simple model of consumer choices under a budget constraint. When a product is not available for a consumer, its price can be thought of as infinite. Therefore, a consumer must optimize without purchasing that product, placing her on a lower indifference curve. When the product is made available, i.e., the price is no longer infinite, the consumer can now maximize utility with all the products available. This chapter attempts to measure how important this change in consumer surplus is relative to price effects.[2] A visual interpretation can be seen in Figure 2.1.

The brewing industry is one such industry where heterogeneous products are offered but products are limited by how much their prices can be changed. As described in detail by Miller and Weinberg (2017) and Weinberg, Sheu, and Miller (2019), brewing firms often priced their products very similarly to their competitors', especially for brands owned by Miller, Coors and Anheuser-Busch Inbev. Several reasons for this similarity were given, such

[2] Alternatively, consumer welfare changes can be modeled within a merger setting where the firm's ability to change prices is limited. When two firms decide to merge, they may decide to change prices or their product offerings. However, this decision must be made with respect to what competitors are doing. For example, prices could rise after the merger due to an increase in market power, but if competition is sufficiently high, this could limit how much prices can increase. In the face of such price constraints, the firm may drop products to increase market power or reduce costs elsewhere. Alternatively, firms may increase product variety by creating new products to capture profit elsewhere if they cannot change prices. Ultimately, any increase in prices and any loss in product variety can lead to consumer welfare loss, and the converse can lead to consumer welfare gains. This scenario is discussed in detail in the next section.

as implicit price collusion, the role of retailers and distributors in price setting and the role of competitors in hampering such implicit collusion through price undercutting. Because of this, firms may decide to cut costs through removing products or capture more profit by introducing new products outside of their flagship brands.

The main data source for this chapter is the IRI Marketing dataset, which provides data on the demand for beer between 2002 to 2012. Using this dataset, I am able to observe in each market which products were removed and which were added before and after the MillerCoors merger of 2008. This merger combined the second and third largest brewery companies in the United States, leading to exogenous changes in costs and market structure. Thus, product varieties changed at the market level, with both product additions and discontinuations after the merger.

Given the evidence of changing product variety from Chapter 1, in this chapter, I design a model of consumer demand to estimate how the measured product variety changes impacted consumers. To make this comparison, I use the model from Miller and Weinberg (2017)'s work, including the same instruments and nesting strategies. I use a random coefficient nested logit model, which is standard in this setting, to model how consumers choose beer: consumers first choose whether they want to purchase a product from the top three competitors and then decide on which specific product to purchase within each nest. I calibrate the model only on the postmerger period to focus on the merger's short-term effects.

After calibration, I estimate the consumer welfare impact of new product additions through a set of counterfactuals. I consider two sets of counterfactuals: one where products newly created after the merger never exist, and one where products discontinued after the merger are never discontinued and remain in their markets. These two counterfactuals estimate the value of losing new products created after the merger and gaining discontinued products after the merger, respectively. The counterfactuals differ in the assumptions on whether firms change prices based on observed product variety, whether the model is estimated individually in each market, the time dimension of model estimation, and whether random coefficients are included. I find that the value of new products created after the merger ranges from 0.05% to 1.35% under these different versions of the model; this is approximately 34% of the consumer welfare effect of coordinated pricing found in Miller and Weinberg (2017). Likewise, I find that the loss in consumer surplus from product discontinuation ranges from -0.14% to -0.175%—approximately -4% of the consumer welfare effect of coordinated pricing found in the same paper.

2.1.1 Literature Review

This chapter contributes to the stream on the effects of changing product variety. I provide a framework for how to model changing product variety, estimating this effect for the merger, and comparing it to the effects of price on consumer welfare.

This chapter contributes to the growing literature on product variety changes and how they relate to consumer welfare. As mentioned previously, Wollmann (2018) studies the impact of product variety changes on consumers under an exogenous shock to product variety through bailouts in 2007. Similarly, Fan and Yang (2022) study how firms may reallocate products after an acquisition of a craft brewery. I make two important contributions complimenting their work. The first is that I focus on a smaller set of markets, avoiding the curse of dimensionality issue that the latter paper addresses. For this reason, I do not need to solve for a large discrete game. This comes at the cost of less granularity of market information. The second contribution is my use of different data. Due to data limitations, the authors of the aforementioned works cannot identify which breweries or beers drive their results. My data are not subject to this limitation, although they cover fewer markets and a shorter time frame. Nevertheless, using these less restricted data allows me to provide more information on firm specifics. By avoiding assumptions on what products leave the market and enter, I provide an empirical test of the conclusions from these authors' models. The central exercise and counterfactual of this chapter can best be compared to the work of Petrin (2002), who studies the value of new products entering the market.

I use tools from the discrete choice estimation literature, notably the random coefficient nested logit model, to estimate demand in this setting. This model, discussed in greater detail in Grigolon and Verboven (2014), combines the random coefficient model and the nested logit model and allows more precise estimates than the standard random coefficient model. This model has been used in a number of other works within this industry and in other settings (e.g., Grennan (2013), Ciliberto and Williams (2014), Conlon and Rao (2020), Miller and Weinberg (2017)). I additionally estimate the model under a standard nested logit, following other works focused on this industry (e.g., Fan and Yang (2022), Hellerstein (2008), Goldberg and Hellerstein (2013), Asker (2016)).

My results both validate and provide additional context to results previously seen in the literature. Examining the merger's effect on product variety, I find results similar to those of Atalay et al. (2023), who find, using a large-scale event study, a slight decline in product variety from a merger. They document the phenomenon of brand consolidation: firms cutting back their product variety in certain brand lines to focus on their highest-revenue products in other brand lines. I observe this phenomenon in my study, as well. After the merger, the number of brands falls within each market, but there is no effect of the merger on the number of products offered per market, implying an offsetting effect coming from new products. Likewise, although Fan and Yang (2022) find a decline in consumer surplus with hypothetical acquisition of smaller breweries, this chapter finds a net increase in consumer welfare from the from the addition of new products and discontinuation of old ones in a merger between larger firms.[3] Therefore, there may be other conflicting factors such as firm size and product substitution that lead to differences in product variety outcomes after a merger.

[3]While this chapter does not examine the supply-side changes occurring due to the merger, evidence of the firms reallocating resources to more efficient processes can be seen in Demirer and Karaduman (2023).

The paper proceeds as follows. In section 2, I describe the setting of the brewing industry, the merger itself, and the suitability of this merger and setting for testing the welfare effects of product variety and price changes. In section 3, I describe the data and provide key summary statistics for each dataset. In section 4, I describe the structural demand models used to estimate the effects of changing product variety on consumer welfare. In section 5, I describe the results of the demand models. In section 6, I describe the main counterfactuals analyzing how consumer surplus changes with changes in product variety postmerger. Finally, in section 7, I conclude.

2.2 Setting

In this section, I expand on the basic intuition on the relationship between product variety and prices and the features of the brewing industry that make it an ideal setting to examine this relationship's effect on consumer welfare. More details on the MillerCoors merger of 2008, the main exogenous shock to market structure that is the focal point of this analysis, can be found in Chapter 1.

2.2.1 Basic Intuition on Relationship Between Product Variety and Prices under a Merger

Here, I provide an expanded intuition on the relationship between product variety changes and price changes to illustrate the claim made in this chapter: that firms trade off prices and quantities. My model uses a simple monopoly profit comparison and shows that, depending on how constrained firms are in changing prices, products may or may not disappear from the market. Due to the large size of the firms considered here, the model provides a baseline for the product variety effects of the merger.

Suppose that I have a market where three firms, denoted A, B and C, operate and sell one good each, $i \in \{1, 2, 3\}$, respectively. The goods are heterogeneous, and prior to the merger, each firm has the following profit condition that determines whether it provides the product in the market:

$$\pi_A^{pre}(c_A) \geq 0 \;\; \& \;\; \pi_B^{pre}(c_B) \geq 0$$

where π_i is the profit for $i \in \{\text{A,B,C}\}$ and c_i represents the costs of production and distribution, which may differ. If profit is greater than zero, the firm provides the product. Firms compete in Bertrand competition.

In this model, a merged firm can change its profits through prices or through costs. Suppose that two of the firms merge and can change only prices. For simplicity, suppose that the third firm's prices remain constant. Denote this new merged firm AB. In the first case, suppose that AB lowers prices. This can occur if the merger affected costs and allowed the merged firm to effectively undercut its competitor. Profits would rise if it is able to capture more customers. In the second case, suppose that AB raises price. This can occur through increased market power from the merger and is more likely if the products are substitutes or one product has a higher cost not reduced by the merger.

Now, suppose that the merged firm can also adjust product varieties. In the first case, suppose that one of the product varieties is dropped. In this setting, this can occur if the products are substitutes and there is some fixed cost of providing the good. Without loss of generality, AB would drop the product if and only if

$$\underbrace{\pi_A^{\text{two products, after}}(c_A) + \pi_B^{\text{two products, after}}(c_B)}_{\text{sum of both product profits}} < \underbrace{\pi_A^{\text{single product}}(c_A)}_{\text{single product profits only}} \tag{2.1}$$

However, suppose that the firm is limited in adjusting prices. In this case, firms may be able to adjust only through changes in product variety. Note that the existence of the third firm and Bertrand pricing allows this to occur naturally within the setting. If the third firm has a low enough price (or marginal cost), price changes may not lead to additional product sales. This can occur if the third firm has low enough costs that it is able to lower prices more than AB can. An alternative reason for constraints on price reductions is coordination. If the firms are coordinating on price, they may want to capture profit elsewhere by lowering costs by removing products with low sales. Therefore, if the firm wants to capture more profit, it could drop a product or introduce a new one in a different market.

I also consider an extension where new products are introduced. After the merger, the firm offers the new product AB if and only if

$$\begin{aligned} &\underbrace{\pi_A^{\text{three products, after}}(c_A) + \pi_B^{\text{three products, after}}(c_B) + \pi_{AB}^{\text{three products, after}}(c_{AB})}_{\text{sum of all product profits}} < \\ &\underbrace{\pi_A^{\text{single product}}(c_A)}_{\text{single product profits only}} \end{aligned} \tag{2.2}$$

and

$$\begin{aligned} &\underbrace{\pi_A^{\text{three products, after}}(c_A) + \pi_B^{\text{three products, after}}(c_B) + \pi_{AB}^{\text{three products, after}}(c_{AB})}_{\text{sum of all product profits}} < \\ &\underbrace{\pi_A^{\text{two products, after}}(c_A) + \pi_B^{\text{two products, after}}(c_B)}_{\text{sum of both product profits}} \end{aligned} \tag{2.3}$$

Note that this would occur with or without limitations in price adjustment, as the main reason that this product exists is the cost changes as a result of the merger.

2.2.2 Implications for Consumer Surplus

To examine what the impacts of these changes on consumer surplus would be, I add an outside good, C, and a small market to this setting. Suppose that there are 9 consumers in this market, with six having preferences for each good as follows: $i > j > k$ for all $i, j, k \in \{A, B, C\}$, without replacement. The last three have preferences only for a good that does not exist prior to the merger, AB, and the outside good C.

I can examine what the final effect on consumer welfare would be given price and product changes depending on the merged firm's choice. If any product's price rises, the consumer welfare of the group of consumers who prefer the product would decline either from their paying a higher price or from their shifting consumption from their preferred good to a cheaper, less preferred option. Alternatively, a price decline would lead to a consumer welfare increase as consumers shift to a more preferred good or pay less for a preferred good. With the inclusion of product variety changes, any withdrawal of a product would lead to a consumer welfare loss or no change at all. This can also be interpreted as prices' being set to infinity—consumers can no longer purchase the good, and for this reason, they may be forced to substitute to another good. For example, suppose that the companies producing A and B merge. This merged company can drop A and force consumers to substitute to B or C. This is a viable strategy if B or C is highly substitutable or if B or C has significant costs that the merged company does not want to undertake.

Likewise, the introduction of a new product in the form of AB can lead to an increase in consumer welfare. When AB is introduced, the three consumers who previously purchased only the outside good may purchase AB, depending on prices. If prices for AB are too high, these consumers will continue purchasing C, leading to no changes in consumer surplus. However, if AB has a low enough price, then individuals will purchase AB, leading to a consumer surplus increase as consumers can now purchase a good that they prefer more at a lower price.

While the effects of price and quantity individually are clear, the combined effects of both are less clear. For example, suppose that prices fall but product variety also falls. Here, consumer surplus would increase from prices but fall from the effects of the product variety changes. Depending on which effect prevails, the merger's effect on product variety would be either net positive or net negative. Evaluating this outcome requires information on both prices and product variety to accurately measure and shut down each channel individually to estimate the consumer welfare effect for each one. For example, in an industry where price changes are constrained, estimating the change in consumer surplus from new and discontinued products postmerger would inform us about the product variety effect of the merger.

This basic intuition underpins the framework of this chapter—a merger's effect on consumer surplus can run through prices or through quantities, and this can have different consumer welfare impacts depending on the merged company's actions. If a firm is constrained in changing price, one option to capture more profit may be to remove or increase product varieties, affecting consumer surplus in the process. This intuition holds with additional firms and additional products as long as price changes are restricted. The following conclusions can be drawn from this observation:

1. If a product is withdrawn from the market, it leads to a weakly negative decline in consumer surplus.

2. If a product enters the market, this leads to a weakly positive increase in consumer surplus.

2.3 Data

This section is the same as Chapter 1 to provide context on the data used for the structural estimation. I additionally include a discussion on the attributes of new and discontinued products used in the structural models. Readers should refer to Chapter 1 for more information on the datasets used in this work.

2.3.1 Consumer Level: IRI Dataset

The IRI dataset provides information on consumer-level demand through scanner data, which show what products consumers buy in stores. I use monthly data from thirty-nine metropolitan statistical areas (MSAs) from 2002 to 2012, which allow me to measure product variety at the final good level and observe revenues, market shares and prices.

The beer industry is a branded consumer product industry, and therefore, varieties, prices and quantities can be measured through supermarket transaction data. The IRI marketing dataset spans 2001–2012 and contains anonymized supermarket transaction data from 51 marketing regions. These marketing regions are typically groups of counties, with some regions crossing state lines. Each observation is an individual sale of a product with a unique UPC identifying a product based on the brand, packaging medium, and size in ounces.

I make several changes to the raw data to facilitate estimation and remove markets with unique legislation restricting the representativeness of observations or impeding my ability to define a market. First, there were several major store-level mergers that occurred in 2001 that affect some of the store-level controls that will prove important for estimation purposes. Therefore, data from the year 2001 are dropped. To better match against state-level data, I do not include data on markets in states that place restrictions on product variety or distribution. Such laws include those limiting the alcohol content of beer sold or prohibiting sales of beer in supermarkets. This criterion leads me to remove eight markets from the dataset. Finally, for estimation purposes and because of unclear definitions of markets, I do not include four markets that consist of entire states. This provides a total of 39 markets in a total of 28 states. For the main demand specification, I additionally subset the brands into the top ten brands by market share. Prior works, such as Miller and Weinberg (2017), Weinberg, Sheu, and Miller (2019), and Ashenfelter, Hosken, and Weinberg (2014), focus on these firms as well. This is partially for computational reasons, but these companies can better be thought of as Miller's and Coors's closest competitors than can local or craft beer brands.

Overall, market shares vary greatly between markets. Figure 2.2 shows the mean market shares of Miller and Coors prior to the merger and of MillerCoors right after the merger and at the end of the sample period. The shares range from less than 3% to nearly 30% for Coors and from less than 1% to over 50% for Miller. The merged company reaches market shares similar to Miller's, ranging from near 0 to over 50% market share.

Figure 2.2: Market Shares for Miller and Coors

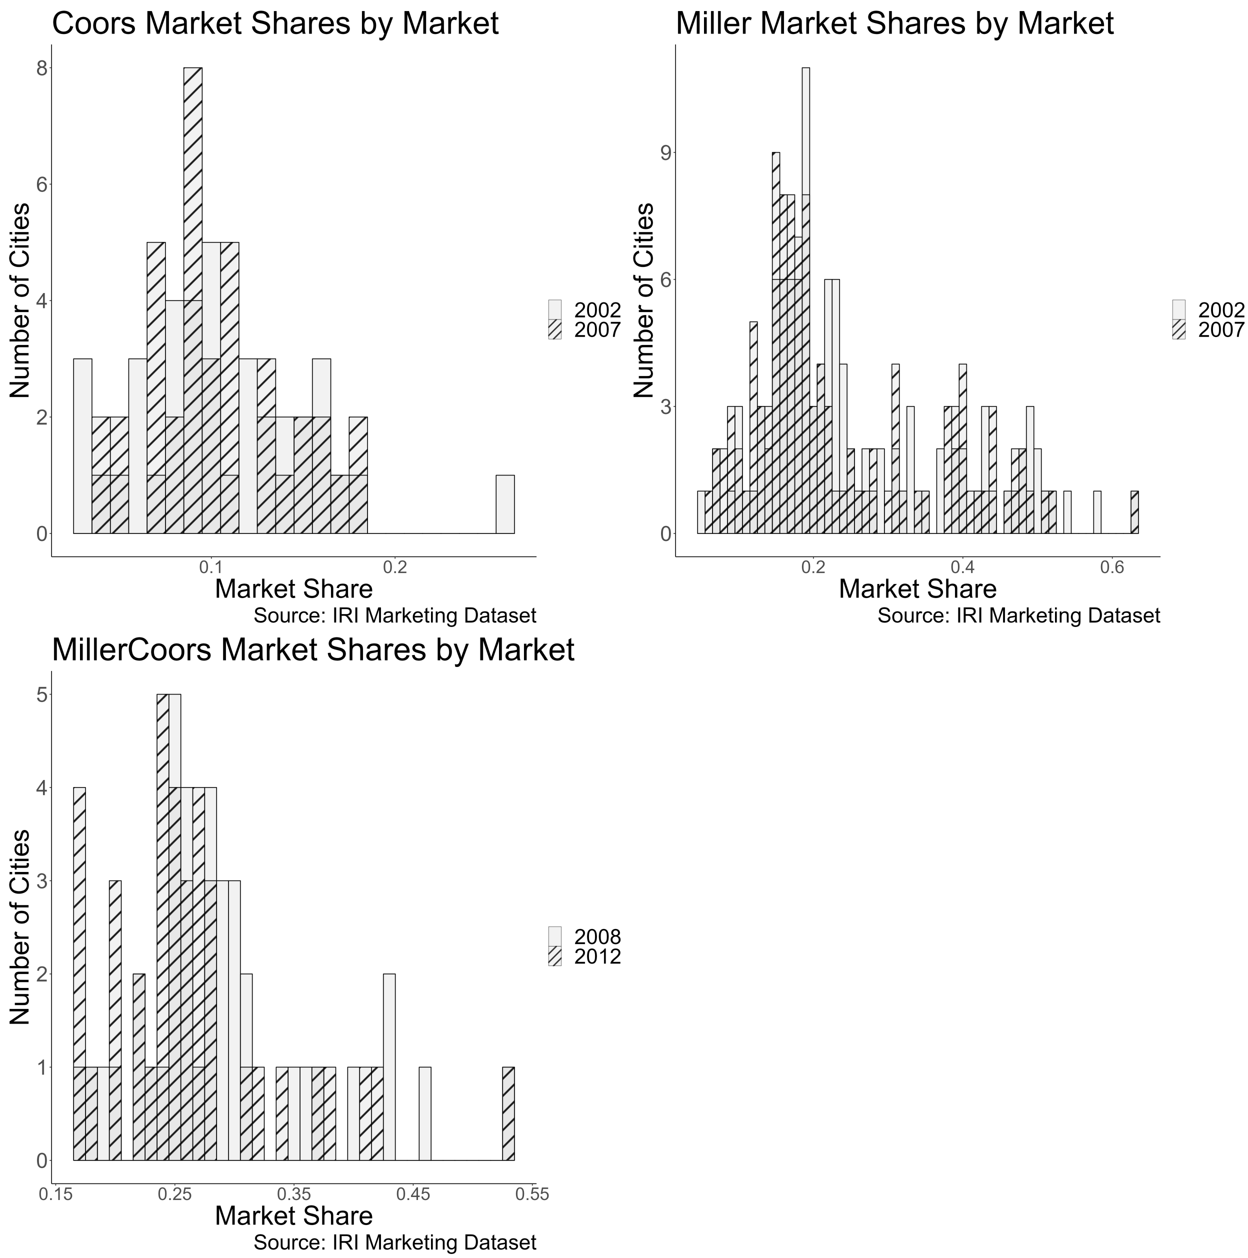

Note: These histograms depict the market share for each market for Coors and Miller for 2002 and 2007 and the market share for MillerCoors for 2008 and 2012. Each observation is a market's annual share for the respective company. Data are compiled through the IRI marketing dataset.

2.3.2 Summary Statistics

I provide the summary statistics for each dataset here for the overall market, for Miller, for Coors and for the combined company MillerCoors for 2002–2012. Of these statistics, I emphasize the measures of distance of the nearest brewery to the nearest market, which the company argued was the main impetus for the merger, and product variety. I show the first main results of the paper, the raw change in product and brand variety, and show that

I first define the difference between product variety and brand variety and examine basic trends to see how the market changed before and after the merger for Miller and Coors. Products in the dataset are defined as a brand × size × packaging type, while brands are names given to products given in the dataset. I consider both definitions for three reasons. First, brands are easily identifiable and clear distinguished in the dataset. For example, Keystone is a different brand from Keystone Light (a lower-calorie version), which is a different brand from Keystone Ice (a version with a higher alcohol by volume). Second, the brand is the highest level of product identification in the dataset beneath the product vendor. If consumers have strong preferences over packaging, such as for twelve packs over twenty-four packs, these results would provide an upper bound on consumer impacts and changes in product variety. Finally, this observable heterogeneity allows me to differentiate between product lines and specific products, which may be important for litigators.

Table 2.1 shows summary statistics for key variables such as revenue, concentration measures and prices. Of key importance for this table are the measures of the number of Coors brands, Miller brands and distance to the breweries for each. Miller had more breweries, leading to an average distance from a Miller brewery to a market of 314 miles. Potentially due in part to this increased capacity, the number of Miller brands is much greater than the number of Coors brands. On average, Miller supplies 27 brands to the entire country, while Coors supplies only 13 brands. The lack of capacity and distance from markets could

Table 2.1: Summary Statistics

Variables	Mean	Standard Deviation	Min	Max
Price of good per product	8.6431	4.6013	0.01	282.5500
Parent company national market share	0.1726	0.1474	0.00	0.3977
Industry HHI	2085.7500	146.2158	1864.01	2304.3722
Parent company regional market share	0.1848	0.1678	0.0000	0.6402
Industry regional HHI	2438.8231	707.3270	1082.2442	4387.0833
Number of Miller products, national	92.6220	27.1676	25.0000	177.0000
Number of Coors products, national	48.2464	12.8936	22.0000	97.0000
Number of MillerCoors products, national	161.8263	35.4402	61.0000	254.0000
Number of Miller brands, national	26.6053	6.6754	10.0000	42.0000
Number of Coors brands, national	12.9904	3.3307	6.0000	23.0000
Number of MillerCoors brands, national	48.9947	8.8273	22.0000	70.0000
Minimum distance from a Miller or Coors brewery, in miles	271.2110	218.0582	12.5304	949.1906
Distance from a Miller brewery, in miles	298.5952	238.1581	12.5304	949.1906
Distance from a Coors brewery, in miles	538.0365	245.8177	77.6465	990.6055
Change in distance from a Miller brewery after merger, in Miles	27.3842	51.1858	0.0000	138.7045
Change in distance from a Coors brewery after merger, in miles	266.8255	244.9850	0.0000	796.9849

Note: This table provides summary statistics for the IRI dataset for Miller and Coors beers and the associated companies between 2002 and 2012. The change in distance is calculated based on the linear distance between a brewery and the centroid of the designated market. The average price is calculated over all beers in every market considered in the study.

potentially affect the cost of producing new products from these sites.

2.3.3 Attributes of New and Discontinued Products Used in the Structural Models

In this section, I describe key attributes of new and discontinued products used in the structural models. Due to the data limitations for model estimation as described in the paper, I can only consider products that are sold in each month or in each quarter per year and products that are sold at least twenty or more times. This limits the number of products considered in the structural estimation. Table 2.2 describes the attributes of new products, compared to the attributes of all new products in the dataset. Table 2.3 describes the attributes of discontinued products, compared to the attributes of all discontinued products in the dataset.

Table 2.2: Attributes of New Products in Dataset and in Structural Subsample

		Mean	Standard Deviation	Minimum	Maximum
New Products					
	Revenue of Products	333023.68	556088.33	2251678.35	10.28
	Total Products Bought	45987.03	65211.1	259939	1
	Market Share	0.0006	0.001	0.0041	1.86×10^{-8}
	Number of Stores Sold at	31.15	29.97	85	1
New Products Used in Structural Model					
	Revenue of Products	679183.81	702753.81	2251678.35	3538.64
	Total Products Bought	77791.42	77131.62	259939	512
	Market Share	0.0012	0.0012	0.0041	6.11×10^{-6}
	Number of Stores Sold at	47.92	33.49	85	6
		Sum			
	Total Revenue of All New Products	19981421.16			
	Total Revenue of All New Products Used in Structural Model	17658779.08			

Note: The table above shows summary statistics between new products used in the structural model versus new products in the dataset.

Table 2.3: Attributes of Discontinued Products in Dataset and in Structural Subsample

		Mean	Standard Deviation	Minimum	Maximum
Discontinued Products					
	Revenue of Products	20295.29	61408.98	284994.95	12.58
	Total Products Bought	8006.19	29331.36	135013	1
	Market Share	3.6434×10^{-5}	0.0001	0.0005	2.1707×10^{-8}
	Number of Stores Sold at	10.38	14.38	57	1
Discontinued Products Used in Structural Model					
	Revenue of Products	15787.18	14568.91	36342.02	3538.64
	Total Products Bought	1925.5	1265.35	3356	512
	Market Share	2.7242×10^{-5}	2.5140×10^{-5}	6.2710×10^{-5}	6.1062×10^{-6}
	Number of Stores Sold at	16.25	17.17	42	7
		Sum			
	Revenue of Products	426201.18			
	Total Revenue of All Discontinued Products Used in Structural Model	63148.76			

Note: The table above shows summary statistics between discontinued products used in the structural model versus discontinued products in the dataset.

Overall, the differences between attributes of new and discontinued products in the structural model subsample directly come from the data restrictions. Overall, new and discontinued products in the model have higher average revenues, bought products, market share and number of stores sold at. Likewise, the standard deviations are lower for all variables due to a smaller sample. Maximums match between the population and the subsample for most variables, and minimums are higher in the subsample. These all are sensible changes due to the restriction made on the data, where these discontinued and new products must be sold more often.

However, the total revenue of new and discontinued products in the population and in the subsample indicate that while most of the new products are captured in the model, not as many discontinued products are. The total revenue of new products in the subsample is \$17,658,779.08, which compared to the population, is \$19,981,421.16, meaning 88% of new product revenue is represented in the model. The total revenue of discontinued products in the subsample is \$63,148.76, which compared to the population, is \$426,201.18, meaning 15% of discontinued product revenue is represented in the model. This is due to the restriction itself. Since a product must be sold at least twenty times, most discontinued products are sold very infrequently in this dataset before ultimately being discontinued. Therefore, while new products are sold and bought almost immediately, discontinued products are eventually phased out, leading to a potential downward bias on the value of discontinued products.

2.4 Structural Model

To fully estimate the effects of changing product variety on consumer welfare, I estimate a model of consumer demand during the merger period. By estimating this model, I can incorporate a flexible system of consumer demand to obtain precise consumer welfare results. Additionally, I can use the model to complete my main counterfactual analysis of what consumer surplus and total surplus would be had the changes in product variety not occurred.

2.4.1 Demand Model

I use the random coefficient nested logit (RCNL) model from Miller and Weinberg (2017) to estimate consumer demand for beer. This model and models analogous to it have been used in a variety of works studying this industry and others. This is my preferred consumer demand model, as it allows me to flexibly estimate consumer preferences for specific types of beer and may better match shopping behavior. Additionally, this allows me to match the price consumer welfare estimates from their models to mine. I use the notation from Miller and Weinberg (2017) to describe an altered version of their model below.

The model is illustrated as follows. Suppose there are $m = 1, \ldots, M$ markets observed during $t = 1, \ldots, T$ time periods. Per market and per time, there are $i = 1, \ldots, N_{mt}$ consumers in each period. Each consumer decides whether to purchase no good (the outside good), or two types of beer: an ale or a lager (defined as a non-ale in the data). Once decided on the type of good, she then chooses a specific good in that category. The products observed are represented as $j = 1, \ldots, J_{mt}$, with the outside good represented as $j = 0$. Products are defined by a combination of brand, the type of good that they are, the packaging type, and the total product size in ounces. Prices are standardized to price per ounce. The conditional indirect utility that consumer i receives from a product j is represented by the following:

$$\begin{aligned} U_{ijmt} = \sigma_1 + (\alpha_1 + \sigma_2 y_i) p_{jmt} + (\alpha_2 + \sigma_3 \nu_1) \mathbb{1}\{\text{Import}\} + (\alpha_3 + \sigma_4 \nu_2) \mathbb{1}\{\text{Light}\} + \\ (\alpha_4 + \sigma_5 \nu_3) \mathbb{1}\{\text{Ale}\} + FE_t + FE_j + \xi_{jmt} + \epsilon^{NL}_{ijmt}(\rho) \end{aligned} \tag{2.4}$$

where y_i is the average income in the market, p_{jmt} is the price of the product, ν_k for $k \in \{1, 2, 3, 4, 5\}$ is the unobserved household shock for the household and the product attributes, FE_t is the year fixed effects, FE_j for product fixed effects, $\mathbb{1}\{\text{Light}\}$ indicates whether the product is a light beer, $\mathbb{1}\{\text{Import}\}$ indicates whether it is an import beer, $\mathbb{1}\{\text{Ale}\}$ indicates whether it is an ale beer, ξ_{jmt} is the product $\times$ market $\times$ time demand shock, and $\epsilon^{NL}_{ijmt}(\rho)$

is the error term, as a function of which type of product was purchased. I normalize the outside good's mean utility to 0, so buyers receive only $\epsilon^{NL}_{ijmt}(\rho)$. Buyers can purchase any lager or ale beer within the dataset, with any beers not in the dataset being part of the outside good.

Under the assumption of a nested logit, I assume the following specification for the error term, given two groups $g \in \{0, 1\}$, where group 1 defines Budweiser, Miller, or Coors products and group 2 other top products. Then,

$$\epsilon^{NL}_{ijmt}(\rho) = \Xi_{igmt} + (1-\rho)\epsilon_{ijmt} \tag{2.5}$$

where ϵ_{ijmt} represents the I.I.D. extreme value draw, Ξ_{igmt} is a draw from a unique distribution such that $\epsilon^{NL}_{ijmt}(\rho)$ is extreme value, and ρ is a nesting parameter between 0 and 1. A larger ρ corresponds to greater correlation for products within the same nest and less substitution between products not in the nest. I also normalize the indirect utility of the outside good such that $U_{i0mt} = \epsilon_{i0mt}$ and assume that the market sizes are the number of unit sales within each region. The outside good contains any beers sold by companies not within the top ten beers of market share, any malt beverages, any other alcohol products such as wine, and beer sold outside supermarkets. This implicitly assumes that these firms in the outside-good group do not price strategically with respect to the firms in this model.

From this specification, I can derive logit choice probabilities for market m and brand j at time t as in Berry, Levinsohn, and Pakes (1995). Multiplying these logit choice probabilities by the market size, I can derive demand as a function of product characteristics, prices, competitor product characteristics and consumer characteristics. The market shares can then be represented as

$$s_{jmt} = \frac{1}{N_{mt}} \sum_{i=1}^{N_{mt}} \frac{\exp(u_{ijmt} - \Xi_{igmt} + (1-\rho)\epsilon_{ijmt})/(1-\rho)}{\exp(I_{igmt}/(1-\rho))} \frac{\exp I_{igmt}}{\exp I_{imt}} \tag{2.6}$$

where I_{igmt} and I_{imt} are the McFadden (1977) inclusive values to normalize the shares. This allows the normalization on the mean indirect utility of the outside good to be $I_{i0mt} = 0$; the inclusive value of the inside products is $I_{igmt} = (1-\rho)\log\sum_{j=1}^{J_{mt}} \exp(u_{ijmt} - \Xi_{igmt} + (1-\rho)\epsilon_{ijmt})/(1-\rho)$ for good type $g \in \{1,2\}$. Finally, the inclusive value for all products is $I_{imt} = \log(1+\exp I_{igmt})$.

Under the assumption that all $\nu_k = 0$, this model reduces to the standard nested logit model, which I use for comparison purposes later. In this case, the utility is linear in parameters:

$$U_{ijmt} = \alpha_1 p_{jmt} + \alpha_2 \mathbb{1}\{\text{Import}\} + \alpha_3 \mathbb{1}\{\text{Light}\} + \alpha_4 \mathbb{1}\{\text{Ale}\} + FE_t + FE_j + \xi_{jmt} + \epsilon^{NL}_{ijmt}(\rho) \quad (2.7)$$

and therefore, the difference in log market shares relative to the mean market share is

$$\begin{aligned} \log(s_{jmt}) - \log(s_{0mt}) = \alpha_1 p_{jmt} + \alpha_2 \mathbb{1}\{\text{Import}\} + \alpha_3 \mathbb{1}\{\text{Light}\} + \\ \alpha_4 \mathbb{1}\{\text{Ale}\} + FE_t + \rho \log(s^{NL}_{jmt|g})\xi_{jmt} \end{aligned} \quad (2.8)$$

where $s^{NL}_{jmt|g} = s_{jmt}/\sum_{j=1}^{J_{mt}} s_{jmt}$ is the conditional share of product j among products within each nest.

Although this model is used for comparison, I prefer using the full RCNL for two reasons. First, the nesting parameter allows flexibility in substitution patterns among similar types of alcohols. Additionally, other papers have adopted the full RCNL model for alcohol-related markets, including the paper whose results I want to directly compare my own with. Second, parameters on taste preferences such as light beer and international brands cannot be fully estimated without random coefficients, and this may be important given a heterogeneous customer base. I provide both the RCNL and the nested logit for comparison.

2.5 Results

In this section, I describe the key results of the paper. The estimation details and instruments, which follow from Miller and Weinberg (2017), can be found in the appendix. I find coefficients consistent with intuition under the demand model and find that new and discontinued products had higher diversion ratios and were more elastic. I then use this model to estimate the two counterfactuals of the paper: what would occur if the new products had not been created after the merger and what would occur if the discontinued products had remained in markets after the merger. I find that the loss in consumer welfare from removing new products far outweighs the gain from including discontinued products.

2.5.1 Demand Model Estimation

To estimate the value of the new and discontinued products, I estimate the demand model after the merger, from 2008 to 2010. I first estimate a standard nested logit for this period of time with monthly and quarterly data. The results are in Table 2.4.

Overall, the aggregate models provide more sensible results than the market-level models. The price coefficients are -0.38 and -0.48 in the monthly and quarterly models, respectively, denoting that for every dollar increase in the price of a twenty-four pack, the probability of purchase falls by 38% and 48%, respectively. Meanwhile, the indicators on imported, ale and light beers are all positive, indicating that the probability of purchase for these products is higher than that for their counterparts. For the market-level models, these are all of the same sign yet much smaller, potentially due to the smaller sample size for each individual model. When examining the elasticities and diversion ratios, I find that the median for the aggregate models is around 4–5 while that for the market-level models ranges from 2 to 3. These are much smaller and imply a less elastic good. Again, this could be due to the lack of data when estimating these models, as there are only a few years of data.

Table 2.4: Demand Logit Results

	Monthly		Quarterly	
	Aggregate Model	Average of Market Models	Aggregate Model	Average of Market Models
	(a)	(b)	(c)	(d)
Price	-0.3813	-0.01086	-0.4837	-0.0098
	(0.0200)	(0.0036)	(0.0492)	(0.0061)
Imported	0.4020		0.63157	
	(0.0794)		(0.0199)	
Ale	0.86153		0.9140	
	(0.0613)		(0.0136)	
Lite	0.4495		0.3675	
	(0.0263)		(0.0638)	
Nested Logit Term	0.0208	0.9515	0.01987	0.9602
	(0.0189)	(0.0150)	(0.0445)	(0.0232)
Observations	155253	155253	65306	65306
Other Statistics:				
All Products:				
Median Own-Price Elasticity	-4.5726	-2.6397	-5.9153	-3.0426
Median Outside Good Diversion Ratio	0.5076	0.0109	0.4917	0.0066
New Products Only:				
Median Own-Price Elasticity	-5.4251	-2.4356	-6.8314	-3.0259
Median Outside Good Diversion Ratio	0.7741	0.1175	0.5096	0.1306

Note: This table shows the results of the nested logit model for the period 2008–2010. The model includes the above variables and the following fixed effects: month, year and indicators on which firm provided the product. For the aggregate model, which calculates demand across all markets and contains all data from 2008 to 2010, I include market fixed effects. For the market-level models, which estimate demand for 2008–2010 for each of the 39 markets, I exclude the characteristic indicators for computational reasons. Only firms included in the original study of Miller and Weinberg (2017) are included. The “pyblp” package is used for model estimation.

Notably, in both models, the diversion ratio for new products is above that for all products, and the elasticities for new products are near or above the elasticities for all products. This result shows that new products are not only more sensitive to price changes but also are more easily substitutable with the outside good. This makes sense, as new products are likely not established enough for consumers to have a strong preference for them. This also shows that any results of the counterfactual may be underestimated: if these products are more easily replaceable, losing them may not decrease consumer welfare as much.

After calibrating the model under a nested logit, I now add complexity to the model and calibrate it under a random coefficient nested logit. For the random coefficient nested logit, I apply the parameters listed in Table 2.5:

Overall, the RCNL's results closely match those of the nested logit model. While the nonrandom coefficients are similar to those under the nested logit model, the random coefficients are all insignificant. Notably, the nested logit terms are also insignificant, like those in the prior model. Both the price elasticities and outside-good diversion ratios are consistent with prior results.

2.6 Counterfactuals

I next examine the following two counterfactuals to examine the value of new products and discontinued products after the merger: namely, I examine what consumer surplus would be if the products newly created after the merger had never been made and what consumer surplus would be if the products discontinued after the merger had continued to exist. From these two counterfactuals, I can compare consumer surplus from the observed baseline and the change in consumer surplus, giving an estimate of the value of these goods. Both of these exercises use the model above, calibrated to the postmerger period. The benefit here is that the products dropped and added after the merger are known, so no assumptions about

Table 2.5: Demand Random Coefficient Nested Logit Results

	Monthly	Quarterly
	(a)	(b)
Price	-0.3813	-0.5580
	(0.0217)	(0.0560)
Price Random Coefficient	0.0000	0.0239
	(0.2765)	(0.0734)
Ale	0.8615	1.0911
	(0.0677)	(0.2723)
Ale Random Coefficient	0.0000	0.0000
	(14.2829)	(13.1826)
Imported	0.4019	0.8515
	(0.1216)	(0.3815)
Imported Random Coefficient	0.0117	-0.0269
	(6.7652)	(16.3632)
Light	0.4496	0.3369
	(0.0352)	(0.0789)
Light Random Coefficient	0.0000	0.0000
	(10.7582)	(10.3504)
Nesting Term	0.0021	0.0027
	(0.0024)	(0.0651)
Random Coefficient Constant	0.0000	0.5036
	(4.9327)	(5.2577)
Observations	155253	65306

	Quarterly	
	MW Firms Only	Top 5% Firms
	(a)	(b)
Other Statistics:		
All Products:		
Median Own-Price Elasticity	-4.5724	-6.6626
Median Outside Good Diversion Ratio	0.5076	0.4654
New Products Only:		
Median Own-Price Elasticity	-5.4248	-7.6794
Median Outside Good Diversion Ratio	0.5478	0.4796

Note: This table shows the results of the random coefficient nested logit model for the period 2008–2010. The model includes the above variables and the following fixed effects: month, year and indicators on which firm provided the product. For the aggregate model, which calculates demand across all markets and contains all data from 2008 to 2010, I include market fixed effects. Only firms included in the original study of Miller and Weinberg (2017) are included. The "pyblp" package was used for the estimation of this model.

which products should be considered in these counterfactuals are needed. The underlying assumption is that every new product was created and every discontinued product dropped because of the merger. If this does not hold, these results provide an upper bound on the value of new products and a lower bound on the value of discontinued products.

An important caveat in these counterfactuals is what happens with prices after any market changes. For this, I consider two cases. In the first, I keep prices constant. This is to consider the scenario where other firms do not respond to any product variety changes. In the second case, I recalculate prices after products are removed from the markets. This is to allow competitors to readjust after seeing their competitor's product variety changes. I calculate these prices through finding the Bertrand price equilibrium after the products have left the market. To estimate these new prices, I assume that the products that will be dropped have a marginal cost equal to 125% of the highest marginal cost estimated in the model. This is to guarantee that these products have a sufficiently large cost that they will not be made.

To provide an appropriate baseline for these results, I replicate as closely as possible the data cleaning done in Miller and Weinberg (2017). Since these consumer surplus results ultimately rely on a specification of demand, the market definition and what data are used to estimate the model, to provide correct relative estimates of consumer surplus changes, I restrict my data such that they mirror the data used to derive Miller and Weinberg's (2017) results as closely as possible. This leads to several major changes regarding discontinued and new products. First, Miller and Weinberg (2017) use three size types for their study: 72-ounce (6-pack), 144-ounce (12-pack), and 288-ounce (24-pack) products. Second, they include products that do not drop out of markets throughout the sample. To replicate this, I consider products available in at least ten markets in each period throughout the sample. Notably, this effects the discontinued product results more, as typically products are discontinued in a few markets before being withdrawn completely. These changes reduce

Table 2.6: Change in Consumer Surplus Estimating Value of New Products

		Time	Δ CS, Prices Fixed	Δ CS, Prices Adjust
Brands added to any market				
	Nested logit	Monthly	1.2503 %	1.1208 %
		Quarterly	1.3511 %	1.2060 %
	RCNL	Monthly	1.2503 %	1.1208 %
		Quarterly	1.3127 %	1.1844 %
New brand allocation for each market				
	Nested logit	Monthly	0.0659%	0.0838%
		Quarterly	0.0508%	0.0585%

the number of new brands studied between 2008 and 2010 and discontinued products studied from 2008 and 2009 from 28 to 10 and from 21 to 4, respectively.

2.6.1 Value of New Products

To estimate the value of new products, I drop a subset of products created postmerger for the period after the merger instead. I consider two subsets: a subset where the only new products are ones added to any market nationally and a subset where, for each market, I consider which products were newly added after the merger. After dropping these products, I re-estimate consumer surplus and estimate the change in consumer surplus under this regime from the baseline of the postmerger period. The final calculation of the value of new products is the difference between the baseline consumer surplus and the consumer surplus under the scenario without the new products, divided by the baseline. I consider two settings where prices do not change and prices readjust as described above. Table 2.6 shows the change in consumer welfare under these four cases for the nested logit.

The results for the counterfactual are consistent across the aggregate models for both the nested logit and RCNL models. In the aggregate models, the value of new products coming from the change in consumer surplus ranges from 1.12% to 1.31%, depending on the time dimension at which the model is estimated and whether the model allows prices to readjust after the change in product variety. In terms of the results of Miller and Weinberg (2017), who estimate a -3.7% decline in consumer surplus due to both the effects of the merger and the effects of coordinated pricing, the effect estimated in the paper is approximately 34% of its magnitude. In terms of Miller and Weinberg's (2017) estimate for unilateral effects only, this model's estimated effect is nearly 60% of the magnitude. Under the market-level models, this effect is much smaller, ranging from at most 0.08% to 0.05%. However, it is still positive, showing that the loss of these new products would still have a negative effect on consumers.

2.6.2 Value of Discontinued Products

To estimate the value of discontinued products, I add a subset of products discontinued postmerger for the period after the merger instead. Due to the decline in the number of brands and therefore products studied, I examine only what happens had the discontinued products been added to their respective markets nationally, rather than on a per-market basis. Additionally, due to the data cleaning, I consider only products discontinued at the end of 2008. The remaining products either did not exist in at least ten markets prior to being discontinued or were not sold in six, twelve or twenty-four packs.

Two assumptions need to be made regarding what the prices and shares would be had the discontinued products existed after 2008. Since these products do not exist in the postmerger period, I have no price data or sense of what their shares would have been. I assume that the prices of the discontinued products follow the same pattern as 2008, the last full year in which the products were available. To calculate the new shares, I start by obtaining the expenditures of these products during 2008 and add them to the total expenditures of all products sold in the market. The shares of the discontinued products from 2008 onward are the expenditures divided by this new total expenditure, and the remaining shares of products in these markets are the expenditures of each product divided by this new total expenditure.

The implicit assumptions here are that the discontinued products' prices changed only with inflation over time and that, with the withdrawal of these products, consumers bought the outside good rather than any products supplied by the companies. This latter assumption biases the results upward, as it is possible that consumers may have substituted to other products marketed by the same company instead of leaving the market entirely.

To do this estimation, I first recalibrate my demand model with data including this new set of discontinued products. The results under this augmented model can be found on Table 2.7. Overall, the results are mostly the same, except that the coefficients on the nested logit parameter are estimated at zero in this model. Since the coefficients in the original model were within two standard deviations of zero in the first nested logit model, this does not seem unusual.

After dropping these products, I re-estimate consumer surplus and estimate the change in consumer surplus under the baseline regime relative to that in the hypothetical situation in which the discontinued products remained on the market in the postmerger period. Here, the counterfactual where prices remain fixed is identical to the outcome seen in the data, where the products were dropped and prices were adjusted by the firms accordingly. The second scenario is one in which prices adjust further after the withdrawal of the products. For example, suppose that a supply-side issue in the counterfactual in which these products existed prevented firms from choosing optimal prices. The final calculation of the value of discontinued products is the difference between the baseline consumer surplus and the consumer surplus under the scenario with the discontinued products, divided by the baseline. Table 2.8 shows the change in consumer welfare under these four cases for the nested logit.

The results for this counterfactual are again consistent across the aggregate models for both the nested logit and RCNL models. In the aggregate models, the change in consumer surplus when discontinued products are retained in their markets ranges from -0.175% to -0.141%, depending on the time dimension at which the model is estimated and whether the

Table 2.7: Demand Random Coefficient Nested Logit Results, Including Discontinued Goods

	Monthly	Quarterly
	(a)	(b)
Price	-0.3925	-0.4973
	(0.0188)	(0.0351)
Imported	0.4263	0.7929
	(0.0745)	(0.1413)
Ale	0.7322	0.4908
	(0.0559)	(0.0794)
Light	0.4244	0.3633
	(0.0211)	(0.0391)
Nested Logit Term	0.0000	0.0000
	(0.0198)	(0.0294)
Observations	166005	76474
Other Statistics:		
All Products:		
Median Own-Price Elasticity	-4.6990	-6.2912
Median Outside Good Diversion Ratio	0.4923	0.4506
Discontinued Products Only:		
Median Own-Price Elasticity	-4.9273	-6.6808
Median Outside Good Diversion Ratio	NA	NA

Note: This table shows the results of the nested logit model for the period 2008–2010 with products that were previously discontinued. The model includes the above variables and the following fixed effects: month, year and indicators for which firm provided the product. For the aggregate model, which calculates demand across all markets and contains all data from 2008 to 2010, I include market fixed effects. Only firms included in the original Miller and Weinberg (2017) study are included. The "pyblp" package is used for model estimation.

Table 2.8: Change in Consumer Surplus Estimating Value of Discontinued Products

		Time	Δ CS, Prices Fixed	Δ CS, Prices Adjust
Brands readded to all markets in which originally discontinued, then withdrawn				
	Nested Logit	Monthly	-0.1412 %	-0.1543 %
		Quarterly	-0.1523 %	-0.1732 %
	RCNL	Monthly	-0.1412 %	-0.1543 %
		Quarterly	-0.1468%	-0.1750%

model allows prices to readjust after the change in product variety. In terms of the results of Miller and Weinberg (2017), who estimate a -3.7% decline in consumer surplus due to both the effects of the merger and the effects of coordinated pricing, the effect estimated in the paper is approximately 4% of that magnitude. In terms of Miller and Weinberg's (2017) estimate for unilateral effects only, this model's estimated effect is 6% of that magnitude.

2.6.3 Summarizing the Two Counterfactuals

These counterfactuals jointly show evidence that the gains from retaining discontinued products are far outweighed by the loss in new product variety after a merger. These results come with several important caveats, which are important for later extensions on this topic. For one, these models do not include a supply side and therefore offer only partial-equilibrium results. Second, my results are benchmarked to those of Miller and Weinberg's (2017) work and should be considered only in the context of their model. Finally, these results provide upper and lower bounds on the value of new products and discontinued products, respectively. The underlying assumption is that all products were discontinued or created because of the merger, which may not necessarily hold. These counterfactuals focus on the most extreme case to provide bounds on what the results could be.

The results are still informative, as they show that product variety changes do matter for consumer surplus effects. New product variety in particular leads to a large gain in consumer surplus. The impact of the removal of new product varieties can be considered of lesser magnitude than that of price coordination, but 34% of the effect may still be important for policymakers. In other industries where product or location variety may be more vital to consumers, such as the drug industry or the supermarket industry, this effect could be significant.

2.7 Conclusion

In this chapter, I examine what the effects of product variety changes postmerger may be for consumers. I first motivate this through a simple model of how the reduction in product variety can weakly negatively affect consumers and then establish the main setting for this chapter: the beer industry and the Miller and Coors merger of 2008. This setting corresponds to a well-established market with a large set of competitors, a merger of two large firms within it, and well-defined product variety that I can easily map to the main data source for this chapter, the IRI marketing database. I then answer what the effects of product variety changes are for consumers by formulating and estimating a model of consumer demand. In the tradition of the empirical industrial organization literature, the model is a random coefficient nested logit model that incorporates product characteristics with the additional flexibility of product nests and idiosyncratic shocks to consumer preferences on product characteristics. I align the model closely to Miller and Weinberg's (2017) work, another paper studying this merger that uses a similar model and the same data. Using my model, I consider what the counterfactuals outcomes would have been had the new products created after the merger not been released exist and had the discontinued products been retained after the merger.

I find that the loss of new products created after the merger would be detrimental to consumers. The change in consumer surplus from the value of new products would be approximately 1.25% or—benchmarked to Miller and Weinberg's (2017) findings—approximately 34% of the change in consumer surplus occurring due to price changes and coordination. Likewise, I find that the loss in consumer surplus from discontinued products would be approximately -0.14% or—again benchmarked to Miller and Weinberg's findings—approximately 4% of the change in consumer surplus occurring due to price changes and coordination. Combining the two, I find that the value added from new products outweighs the value lost from discontinued products.

Future work can incorporate a supply-side model to find the general equilibrium and incorporate more dynamics of how product variety can change on the producer side. Work such as Wollmann (2018) provides a baseline for how to incorporate how the firm changes product variety based on the fixed costs of the product. Ultimately, this chapter also serves as a guideline for how product variety changes should be studied for mergers. Works such as Atalay et al. (2023) emphasize the increasing importance of these effects, and policymakers could use this study's framework to model how they study changes in product variety for a merger.

2.8 Appendix

2.8.1 Nested Logit Model Estimation

In this section I describe the estimation procedure in detail and which instruments are used. These processes are identical to the procedures used in Miller and Weinberg (2017). The code used to create this version model uses the pyBLP package in python developed by Conlon and Gortmaker (2020), which incorporates multiple best practices in estimation. I direct readers to Conlon and Gortmaker (2020) and Miller and Weinberg (2017) for further details on the package itself and additional details on the estimation strategy, respectively.

Estimation Strategy

Central to the estimation strategy is the nested fixed point procedure of Berry, Levinsohn, and Pakes (1995). I estimate the model given the population moment condition $E[Z'\omega(\theta_0^D)] = 0$, where $\omega(\cdot)$ is defined as per the optimization solution, $\theta_0^D = (\alpha, \Sigma, \rho)$ are the population parameters, where Σ is a diagonal matrix consisting of the σ parameters on the random coefficient shocks, and α consists of all α parameters on the model, and Z is a matrix of instruments. This model is solved through a generalized method of moments (GMM) estimator where I solve for the following for some positive definite weighting matrix A:

$$\theta^D = \arg\min_{\theta} \omega(\theta)' Z A^{-1} Z' \omega(\theta) \tag{2.9}$$

For any set of candidate parameters, this contraction mapping identifies the mean utility that matches observed and predicted market shares.

I use the two-step Hansen procedure for GMM estimation. In the fist step, $A = Z'Z$. Once the error is minimized, I re-estimate the model using an optimal weighting matrix that uses cluster correction to correct for autocorrelation, heteroscedasticity, and within-state correlations. As the number of regions increases, asymptomatic consistency is reached.

Instruments

In order to estimate the model, an instrument for price, an instrument for the nests, and an instrument for each nonlinear parameter is needed. Since prices are likely to be correlated with the structural error term given firms price optimally given product- and market-specific consumer values, I have a standard endogeneity issue. Additionally, due to the presence of the nests and the nonlinear parameters, I have a simultaneity issue where demand parameters cannot be jointly separated from market shares. Therefore, I need to address the endogeneity from these demand parameters and the nesting parameter, which helps us separately estimate the two.

To best compare the model from Miller and Weinberg's 2017 model, I use the instruments from their replication files. This is to best compare their model's findings to this model and keep as much similarity between the two models as possible. This provides the following instruments: A group indicator for whether the product is either Miller or Coors at or after 2008, distance from the nearest brewery times the price of gas for both these breweries and the other breweries in the dataset, the number of products in the market for both groups, the distance from breweries to each market, summed, these two variables interacted with dummies for whether the product is MillerCoors or Anheuser-Busch InBev, and the mean income of each market interacted with the imported, lite and ale dummies. Justification for these instruments can be found in their paper. One key caveat of using their instrument requires an assumption for the instrument that identifies the nested logit parameter, which is the correlation between unobserved preferences for goods within the nest. For this, an instrument is needed such that it is uncorrelated with the structural error term, and provides exogenous variation in the conditional share of the goods within the nest. For this, I use the number of products in each nest as in Miller and Weinberg (2017). Although the main counterfactuals rely on product variety, I justify this through a timing assumption. I assume that firms choose product varieties per year, then, conditional on observables (characteristics of the good) and unobservables (random shocks), consumers then choose which product type they want and what specific product.This argument is akin to the selection on observables and unobservables argument made in Dale and Krueger (2002). The number of products in each nest is a standard instrument that is negatively correlated with the conditional share. Finally, I need instruments for parameters on consumer heterogeneity for preferences on characteristics. For this, I use mean income interacted with each product characteristic.

These characteristics are different than Miller and Weinberg (2017). I assume that the structural error term is mean independent of income and product characteristics, a standard assumption in the literature. This provides me with a total of 12 instruments.

2.8.2 Alternative Counterfactuals

In this section, I describe the two alternate counterfactuals conducted to estimate what potential effect the merger could have on product variety. Two estimate the value of new products when discontinued products also exist in the market, and the value of discontinued products if no new products existed in the market. These counterfactuals require the second model which includes discontinued products and new products together, so all assumptions made about discontinued products existence in the market are kept. More details can be found within the main paper on these assumptions.

Value of New Products if Discontinued Products Remained in the Market

In this counterfactual, I estimate the value of new products under the scenario where discontinued products remain in the market the entire time. This counterfactual considers a scenario where if the merger were to not occur, both new products would not be created and discontinued products would not be eliminated. This model puts stronger assumptions on how a firm may act after the merger by stating the merged firm would keep all products in the market and not remove discontinued products as observed in the current equilibrium.

To estimate this, I first estimate consumer welfare in a version of the model where the underlying data includes new and discontinued products. I remove new products from the data, and re-estimate consumer welfare. Taking the difference and dividing it by consumer welfare under the case where both new and discontinued products were in the market, I am able to provide an estimate for the value of new products under this alternate scenario. I again consider two cases: one where prices remain static, and another where prices readjust after observing the change in product variety. Table 2.9 shows the results for this counterfactual.

Table 2.9: Change in Consumer Surplus Estimating Value of New Products if Discontinued Products Remain

		Time	Δ CS, Prices Fixed	Δ CS, Prices Adjust
Brands added to any market				
	Nested Logit	Monthly	1.0491 %	0.9585 %
		Quarterly	0.7099 %	0.6409 %
	RCNL	Monthly	1.0491 %	0.9584 %
		Quarterly	0.7087 %	0.6401 %

Overall, I find the value of new products is slightly lower than the baseline results in the paper, but still positive. The estimates range from 0.64% to 1.04%, which for comparison, the results for the baseline model range from 1.12% to 1.35%. One potential explanation for this could be due to substitution between new and discontinued products. If new and discontinued products were substitutes of each other, then keeping both in the market would dampen the effect of new products since discontinued products already exist and serve similar customers. Examining this issue further would require both products to exist within the market concurrently, which does not hold for a long enough period of time within my dataset.

Value of Discontinued Products if New Products Never Existed

In this counterfactual, I estimate the value of discontinued products under the scenario where new products were never created. This counterfactual considers a scenario where if the merger were to not occur, both discontinued products would not be removed and new products would not be created at all. This model puts stronger assumptions on how a firm may act after the merger by stating the merged firm would still discontinue products and not create new products if a merger were to occur.

To estimate this, I first estimate consumer welfare in a version of the model where the underlying data includes no new products but discontinued products. I then remove discontinued products from the data, and re-estimate consumer welfare. Taking the difference and dividing it by consumer welfare under the case where there were no new products but discontinued products were in the market, I am able to provide a change in consumer welfare value for discontinued products. I again consider two cases: one where prices remain static, and another where prices re-adjust after observing the change in product variety. Table 2.10 shows the results for this counterfactual.

Table 2.10: Change in Consumer Surplus Estimating Value of Discontinued Products if New Products Never Existed

		Time	Δ CS, Prices Fixed	Δ CS, Prices Adjust
Brands readded to all markets in which originally discontinued, then withdrawn				
	Nested Logit	Monthly	-0.0351%	-0.0369%
		Quarterly	-0.0331%	-0.0353%
	RCNL	Monthly	-0.0351%	-0.0369%
		Quarterly	-0.0331%	-0.0353%

Overall, I find the value of discontinued products is lower than the baseline results in the paper, but still results in a consumer welfare loss. The estimates range from -0.03% to approximately -0.04%, which for comparison, the results for the baseline model range from -0.14% to -0.17%. One potential explanation for this could be consumers may substitute towards other goods in the market or go to the outside good more often than if there were no new products. Discontinued products could matter less when no new products from the merged firm exist to potentially take their place.

2.9 References

Abadie, Alberto, and Javier Gardeazabal. 2003. "The Economic Costs of Conflict: A Case Study of the Basque Country." *American Economic Review* 93 (March): 113–32. https://doi.org/10.1257/000282803321455188.

Ashenfelter, Orley C., Daniel S. Hosken, and Matthew C. Weinberg. 2015. "Efficiencies Brewed: Pricing and Consolidation in the US Beer Industry." *The RAND Journal of Economics* 46 (June): 328–61. https://doi.org/10.1111/1756-2171.12092.

Ashenfelter, Orley, Daniel Hosken, and Matthew Weinberg. 2014. "Did Robert Bork Understate the Competitive Impact of Mergers? Evidence from Consummated Mergers." *Https://Doi.org/10.1086/675862* 57 (August): S67–100. https://doi.org/10.1086/675862.

Asker, John. 2016 "DIAGNOSING FORECLOSURE DUE TO EXCLUSIVE DEALING*." https://doi.org/10.1111/joie.12114.

Association, Brewers. 2022. "National Beer Sales & Production Data." https://www.brewersassociation.org/statistics-and-data/national-beer-stats/.

Atalay, Enghin, Alan T. Sorensen, Christopher J. Sullivan, and Wanjia Zhu. 2023. "Product Repositioning by Merging Firms," May. https://doi.org/10.3386/W31229.

Barajas, Jesus M, Geoff Boeing, and Julie Wartell. 2017. *Untapped: Exploring the Cultural Dimensions of Craft Beer.*

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica* 63 (July): 841. https://doi.org/10.2307/2171802.

Capobianco, Antonio. 2018. "Non-Price Effects of Mergers." *129th OECD Competition Committee Meeting*. https://www.oecd.org/competition/non-price-effects-of-mergers.htm.

Chandra, Ambarish, and Matthew Weinberg. 2018. "How Does Advertising Depend on Competition? Evidence from u.s. Brewing." *Https://Doi.org/10.1287/Mnsc.2017.2889* 64 (January): 5132–48. https://doi.org/10.1287/MNSC.2017.2889.

Ciliberto, Federico, and Jonathan W Williams. 2014. "Does Multimarket Contact Facilitate Tacit Collusion? Inference on Conduct Parameters in the Airline Industry." *Journal of*

Economics 45: 764–91. https://doi.org/10.1111/1756-2171.12070.

Conlon, Christopher, and Jeff Gortmaker. 2020. "Best Practices for Differentiated Products Demand Estimation with PyBLP." *RAND Journal of Economics* 51 (December): 1108–61. https://doi.org/10.1111/1756-2171.12352.

Conlon, Christopher T., and Nirupama L. Rao. 2020. "Discrete Prices and the Incidence and Efficiency of Excise Taxes."" American Economic Journal: Economic Policy 12 4 111-43 2020 10.1257/pol.20160391 https://www.aeaweb.org/articles?id=10.1257/pol.20160391

Dale, Stacy Berg, and Alan B. Krueger. 2002. "Estimating the Payoff to Attending a More Selective College: An Application of Selection on Observables and Unobservables." *The Quarterly Journal of Economics* 117 (November): 1491–1527. https://doi.org/10.1162/003355302320935089.

Demirer, Mert, and Omer Karaduman. 2023. "Do Mergers and Acquisitions Improve Efficiency: Evidence from Power Plants," January. https://gsb-faculty.stanford.edu/omer-karaduman/files/2022/12/Draft.pdf.

Division, Antitrust. 2008. "STATEMENT OF THE DEPARTMENT OF JUSTICE's ANTITRUST DIVISION ON ITS DECISION TO CLOSE ITS INVESTIGATION OF THE JOINT VENTURE BETWEEN SABMILLER PLC AND MOLSON COORS BREWING COMPANY."

Elzinga, Kenneth G., Carol Horton Tremblay, and Victor J. Tremblay. 2015. "Craft Beer in the United States: History, Numbers, and Geography*." *Journal of Wine Economics* 10 (December): 242–74. https://doi.org/10.1017/JWE.2015.22.

Fan, Ying, and Chenyu Yang. 2022. "Estimating Discrete Games with Many Firms and Many Decisions: An Application to Merger and Product Variety." *SSRN Electronic Journal*, May. https://doi.org/10.2139/SSRN.3681556.

Goldberg, Pinelopi Koujianou, and Rebecca Hellerstein. 2013. "A Structural Approach to Identifying the Sources of Local Currency Price Stability." *The Review of Economic Studies* 80 (January): 175–210. https://doi.org/10.1093/RESTUD/RDS015.

Goodman-Bacon, Andrew. 2021. "Difference-in-Differences with Variation in Treatment Timing." *Journal of Econometrics* 225 (December): 254–77. https://doi.org/10.1016/J.JECONOM.2021.03.014.

Grennan, Matthew. 2013. "Price Discrimination and Bargaining: Empirical Evidence from Medical Devices." *American Economic Review* 103 (February): 145–77. https://doi.org/10.1257/AER.103.1.145.

Grigolon, Laura, and Frank Verboven. 2014. "Nested Logit or Random Coefficients Logit? A Comparison of Alternative Discrete Choice Models of Product Differentiation." *The Review of Economics and Statistics* 96 (December): 916–35. https://doi.org/10.1162/REST_A_00420.

Hellerstein, Rebecca. 2008. "Who Bears the Cost of a Change in the Exchange Rate? Pass-Through Accounting for the Case of Beer." *Journal of International Economics* 76 (September): 14–32. https://doi.org/10.1016/J.JINTECO.2008.03.007.

Heyer, Ken, Carl Shapiro, Jeffrey Wilder, Ken Heyer, Carl Shapiro, and Jeffrey Wilder. 2009. "The Year in Review: Economics at the Antitrust Division, 2008-2009." *Review of Industrial Organization* 35 (December): 349–67. https://doi.org/10.1007/S11151-009-9232-1.

Khmelnitskaya, Ekaterina, Guillermo Marshall, and Scott Orr, 2024. "Identifying Scale and Scope Economies using Demand-Side Data". *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4757085

Martin, Andrew. 2007. "Merger for SABMiller and Molson Coors." *New York Times*. https://www.nytimes.com/2007/10/10/business/worldbusiness/10beer.html.

Mcfadden, Daniel. 1977. “Modelling the Choice of Residential Location.” https://elischolar.library.yale.edu/cgi/viewcontent.cgi?article=1709&context=cowles-discussion-paper-series.

Miller, Nathan H., and Matthew C. Weinberg. 2017. “Understanding the Price Effects of the MillerCoors Joint Venture.” *Econometrica* 85 (November): 1763–91. https://doi.org/10.3982/ECTA13333.

Petrin, Amil. 2002. “Quantifying the Benefits of New Products: The Case of the Minivan.” *Journal of Political Economy*. Vol. 110.

Tremblay, Victor J., Carol Horton Tremblay, Victor J. Tremblay, and Carol Horton Tremblay. 2005. “The US Brewing Industry: Data and Economic Analysis” 1: 379. https://EconPapers.repec.org/RePEc:mtp:titles:0262201518.

Weinberg, Matthew, Gloria Sheu, and Nathan Miller. 2019. “Oligopolistic Price Leadership and Mergers: An Empirical Model of the u.s. Beer Industry.” *2019 Meeting Papers*. https://ideas.repec.org/p/red/sed019/1210.html.

Wollmann, Thomas G. 2018. “Trucks Without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles.” *American Economic Review* 108 (June): 1364–1406. https://doi.org/10.1257/AER.20160863.

Chapter 3

The Effects of Domestic Restrictions in Government Procurement for Transportation Goods

3.1 Introduction

Government procurement can be costly for transit goods, yet is vital for constituents. In the United States, total Federal Transportation Agency (FTA) grants for bus and rail transit procurement exceeded $ 3.6 billion dollars in 2018[1]. Despite the high cost, one in ten Americans use public transit on a daily or weekly basis[2], making this a vital good for many individuals. Therefore, policies surrounding rail and bus purchases via procurement are of high importance to bringing these goods to individuals at the lowest cost to the government. Additionally, these transit procurement contracts may have positive spillovers to the labor market, especially for steel production (Platzer et al., 2019). Due to these large costs,

[1]See FTA Statistical summaries here: https://www.transit.dot.gov/funding/grants/statistical-summaries

[2]See Pew Research, https://www.pewresearch.org/fact-tank/2016/04/07/who-relies-on-public-transit-in-the-u-s/

local transit agencies typically rely on federal funds for procurement. To guarantee positive spillovers, or for other political economy reasons, federal funds often come with requirements. One such requirement is Buy America (BA hereafter), which requires procurement purchases using federal funds for buses and railcars to reach some threshold of domestic content.[3].

BA underwent significant changes in 2015 that potentially could raise government procurement costs. Until 2017, BA required procurement purchases using federal funds for buses and railcars to contain at least 60% domestic content. With the introduction of the Fixing America's Surface Transportation Act (FAST) Act in 2015, this domestic content requirement increased to 65% for Fiscal Year 2018 to Fiscal Year 2019 and then to 70 % for Fiscal Year 2020 onwards. Recent changes in the BA bylaws due to the FAST Act pose several questions. The primary question is what the cost or benefits of the policy change are. This policy could lead to less competitive procurement by eliminating competitive foreign competitors, leading to increased costs for agencies. However, the policy is intended to increase the share of procurement going to domestic manufacturers. This could imply incumbent domestic producers winning more contracts, foreign companies investing more in domestic manufacturers, or a combination of the two, all benefiting domestic firms. Additionally, this policy has greater implications for trade policy concerning foreign involvement in high cost domestic products and in public finance through considering how changes in the cost burden of government procurement can increase or decrease economic welfare by reducing the need for distortionary taxation.

This chapter investigates how restrictions on firms participating in government procurement auctions affect local government procurement costs, firm participation within procurement and firm investment decisions. In this chapter I focus on the "Buy America" bylaws. This plausibly exogenous policy change in 2015 allows me to examine the effects of the policy before and after, and allows me to examine how firms and governments reacted to the policy.

[3]The U.S. Department of Transportation website has a comparison of the BA Provisions here.

I propose three structural models to study the policy's effect on procurement decisions and show how the effects of BA can be ambiguous in two of these models.

An illustrative example of the policy can be found below in Figure 3.1 and Figure 3.2. Consider a buyer choosing between a domestic and a foreign firm for procurement. To win the auction, the winning firm must have the lowest cost. In Figure 3.1, foreign firms are assumed to have a normal distribution with mean C and domestic firms are assumed to have a normal distribution for cost with mean $C' > C$. Compared to the average cost of domestic firms, the foreign firm is expected to win more procurement contracts. In contrast, Figure 3.2 shows domestic restrictions enacted on the foreign firm. Given that the foreign firm acted optimally before the policy, the mean cost for the foreign firm should shift by a factor $\epsilon \geq 0$. The effect of the policy then depends on how large the cost distribution shifted by. If ϵ is sufficiently large as in Figure 3.2, domestic firms could have lower costs on average and and thus win procurement contracts more often.

Figure 3.1: Procurement Distribution Before Policy

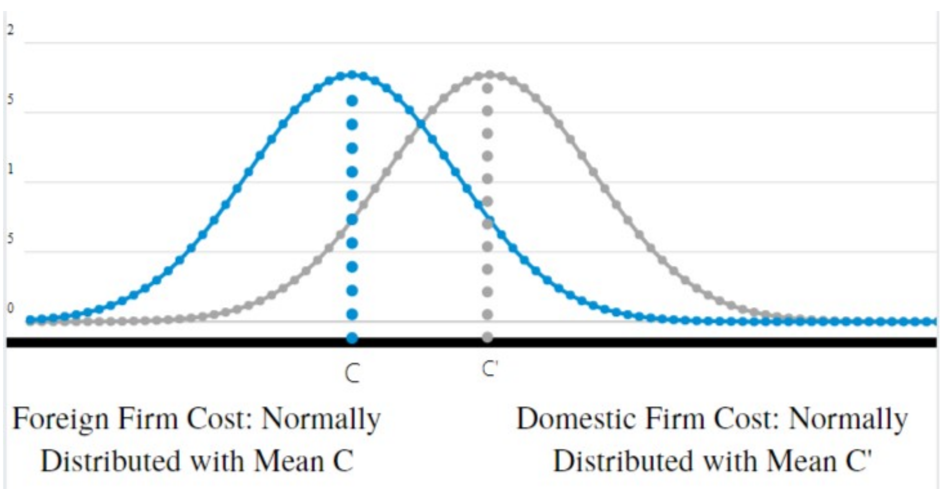

An alternative model could view Buy America as a fiscal cost imposed to participate in procurement auctions for foreign firms. In this case, suppose a foreign firm invests in their domestic holdings and pays a cost to domestic manufacturers to do so. The foreign firm is

Figure 3.2: Procurement Distribution After Policy

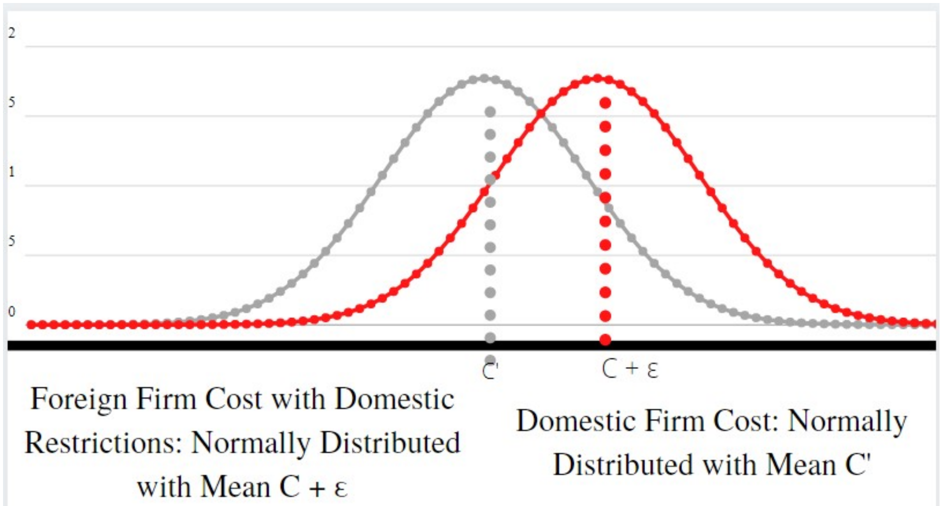

then no longer subject to the penalty and acts as if they were a domestic firm. If this is a goal for policymakers, the higher the investment costs, the more foreign firms must invest to participate, increasing the welfare of domestic suppliers.

The illustrative example provides one such framework for examining the effects of BA. The goal is to measure cost distributions using an auction framework before and after the policy. Under the provided example, BA could lead to higher procurement costs due to the higher costs of domestic goods compared to international goods. If procurement costs stay the same or decreased, then this policy may have not had the desired effect lawmakers wanted.[4] This could imply either weak enforcement, a transition to a different system or payment for procurement or some firm-level decisions in the wake of the policy. Alternatively, the goal of the policy could be to induce more foreign companies to invest in domestic manufacturers and become domestic firms themselves. Costs here could fall as more domestic competitors compete for contracts. Finally, the role of market power may play a role. Given that foreign firms are penalized, domestic firms could increase their bids to gain additional

[4]These restrictions on government contracts could also have a national security purpose. While this chapter acknowledges this concern, most procurement of municipal buses and railcars comes from companies originating from Canada, Mexico and Japan. Additionally, some of these firms already have factories within the United States, and this policy would increase investment in those domestic facilities.

profit. This can lead to greater procurement costs.

To test whether these changes had any effect, I first examine aggregate cost changes from the Federal Transit Authority (the sole provider of grants for transportation goods to local agencies and the agency where BA applies). I find two potential stories depending on which good is being examined. For the bus industry, the policy seems to have negligible effects on procurement costs after 2017. For the railcar industry, the policy seems to have a major effect on procurement costs after 2017. Anecdotal evidence suggests there was investment and adjustment in the bus industry, while the railcar industry was unable to adjust. This provides justification for creating structural models examining this behavior.

I create three models based on existing auction models and discussing identification and testability in these frameworks. The first, the standard auction model, relies on a Guierre, Perrigne, and Vong (2000) (GPV (2000) hereafter) estimation framework and starts as a baseline comparison. The second incorporates heterogeneity between firms and uses Krasnokutskaya and Seim (2011)'s framework of penalized auction types. The final model nests the second model and includes endogenous entry given a noisy signal from Roberts and Sweeting (2013). I include the novel extension of allowing a firm to switch its type in an auction. This addition allows policymakers and economists to analyze investment costs and identify how this policy induced more foreign firms to invest in domestic producers. I then discuss possible extensions and shortcomings of these models.

The paper is divided into six sections. In the second section, I discuss the literature on auctions and government procurement. In the third section, I give background on BA and the FAST Act. I discuss the assumptions, identification strategies, and testability for each of the three models in the fourth section. In the fifth section, I walk through an empirical exercise using the three models. I conclude with the fifth section. The Appendix discusses the empirical framework which I plan on using to answer this research question empirically.

3.2 Literature Review

This chapter contributes to literature on government procurement, domestic content policies, and estimating auctions. Three papers are central to this work: GPV (2000), Krasnokutskaya and Seim (2011) and Roberts and Sweeting (2013). Each paper provides a framework for each of the three models described later to examine the effects of Buy America. In particular, GPV (2000) estimates a standard first price auction model using an indirect method. Krasnokutskaya and Seim (2011) study preferential treatment for small-business in procurement. To do so, they consider an auction model with heterogeneous types of bidders where one group is penalized. Finally, Roberts and Sweeting (2013) study whether auctions are optimal in dynamic settings versus a sequential mechanism. Of note is their model of endogenous participation for auctions where each player observes a noisy signal of their value before participation. I build upon these three elements of their models and incorporate it into my own. I also allow the bidders to change their bidder type before deciding on participation. For identification and testability, Athey and Haile (2002) and Gentry and Li (2014) provide frameworks for the standard model and the endogenous entry model. Specifically, Athey and Haile (2002) consider identification and testability in various types of auction models, including first price standard auctions and auctions with heterogeneous types. They also consider identification and estimation with limited data, which is important for application. Gentry and Li (2014) do similar exercises for an endogenous entry model where players receive a noisy signal for their valuation before participation. These papers provide the econometric framework for my models.

This chapter provides a new framework by which to study government procurement. Work by Lewis and Bajari (2011) study time-based incentives in a scoring auction model in the context of Caltrans road improvements. Similar to this chapter, the authors consider an endogenous participation model to answer this question. Bajari, Houghton and Tadelis

(2014) study government procurement in relation to scoring auctions and government costs. While I also plan on examining the costs of the BA policy, my models show the policy may have an ambiguous effect on government cost. Additionally, I plan on using the simpler framework of a First Price Auction instead of a scoring auction.[5] Other papers such as Li and Zheng (2009) study entry behavior in a first price auction. Of note is their pure strategy entry model where potential bidders draw private costs before entry. I consider a variant of this model where private costs are noisy. Jehiel and Lamy (2015) consider discrimination against foreign bidders in an endogenous entry model. I expand the model to consider an empirical exercise from policy.

Finally, this chapter expands on theoretical arguments made for domestic content limits. Researchers previously studied a similar act, the Buy American Act, in theoretical settings. Miyagiwa (1991) studies whether discriminatory procurement practices such as Buy American reduces imports. Similarly, Naegelen and Mougeot (1998) examine Buy American in a first-price auction framework. They find that a variant of a First Price Sealed Auction can be the optimal mechanism in this scenario. My model relaxes some of these assumptions on government preferences and the players. One of the more related theoretical papers to my work is Arozamena and Cantillon (2004), which examines investment incentives in procurement auctions. In their framework, they find implications of auction models where investment decisions are made before the auction. My model considers a special case of their model, where the bidding functions, investment distributions, and players are specified.

3.3 Setting

In this section, I discuss the Buy America Policy and the changes that occurred with the 2015 FAST Act. I then provide a discussion of the short run impacts of the policy, which

[5]The majority of auctions in the data set are completed via an invitation for bids (IFB), which is similar in structure to a first price auction. Therefore, I do not discuss Requests for Proposals in this paper, although that is a potential avenue for future work.

motivates the later auction models.

3.3.1 Buy America

In 1978, Congress passed the Surface Transportation Assistance Act containing the first domestic content law for transit goods. This mandate, later known as Buy America, determines how federal funds can be used in procurement of transportation goods (Platzer and Mallett, 2019). From 1978 to 2015, few revisions of the law have been made.[6] Central to the law is the definition of "domestic good". The definition of domestic depends on what type of good is being purchased. For iron, steel and manufactured products, this domestic content limit is 100%, meaning it must be completely manufactured and built in the United States. For buses and rolling stock (including train control, communication, and traction power equipment), this domestic content limit is 60%. Labor costs are not included in this domestic content limit and all final production must occur within the United States.[7]

Firms can waive BA under three main exceptions.[8] The first is for public interest reasons. The second is if the goods needed for procurement are produced at an insufficient quality or quantity. The third is if the purchase of a domestic good leads to a cost increase of more than 25% as compared to a similar foreign good (49 CFR Chapter VI Part 661.7, 2010). This exception is a crucial parameter for examining the impact of BA, as it provides an upper bound on domestic costs compared to foreign costs.

If these exemptions do not apply and the domestic content percentage is 60% or more,

[6]The list of historical changes to the law can be seen on the "Regulations and Guidance" page on the FTA website: `https://www.transit.dot.gov/regulations-and-guidance/buy-america/buy-america-regulations`

[7]While the FTA is not the only organization that distributes funds, it is the least binding policy in terms of the domestic content limit. Other government agencies have a full 100% domestic content limit for their goods. Additionally, the FTA is the only organization that was affected by the FAST Act. The full table of comparsons between various BA type policies can be seen on the FTA website: `https://www.transportation.gov/sites/dot.gov/files/docs/buy_america_provisions_side_by_side.pdf`

[8]All Waivers Listed can be seen on the FTA website `https://www.transit.dot.gov/regulations-and-guidance/buy-america/waivers-granted`

then the good is BA compliant and the local agency can use FTA Federal funds for purchasing this good. For further verification, pre- and post- purchase audits must be completed for rolling stock. BA is largely binding for large purchases, as a General Public Interest Waiver exists for exceptions of purchases below $150,000.

3.3.2 Fixing America's Surface Transportation Act (FAST Act)

The FAST Act, signed into law in December 2015, provided funding for surface transportation infrastructure and investment (FTA, 2019). While the majority of the act focused on motor transportation, the FAST Act also changed the definition of Domestic for rolling stock and buses beginning in fiscal year (FY) 2018. The definition of domestic for BA purposes regarding rolling stock and buses was increased to 65% from FY 2018 to FY 2019, then to 70% for FY 2020 onward. In addition, waivers required additional information such as certification that the foreign steel, iron and manufactured goods met quality and quantity standards.

Some anecdotal evidence supports the argument that the domestic content limit change was a plausibly exogenous shock to firms and local governments. The American Public Transportation Association, a nonprofit advocacy group of public and private sector transportation organizations, published a notice in May 2016 disagreeing with the domestic content increases proposed in the FAST Act in the prior year. Additionally, they noted "Pre- and post-delivery audits, under this proposal, would take on a new complexity and likely a new cost" (APTA, 2016). These costs would be undertaken by local transit agencies.

Further comments under the Federal Register show additional uncertainty with policy implementation. For example, agencies expressed confusion about previously scheduled procurement decisions. The official guidelines state "for purchase orders placed against State purchasing schedules before October 1, 2015, for the delivery of rolling stock in FY 2018 and beyond, the increased domestic content requirements will not apply. For purchase orders placed against State schedules on or after October 1, 2015, for rolling stock that will be delivered in FY 2016 and 2017, the domestic content requirement must exceed 60%" (Federal Register, 2016). Regardless of when the contract was set, by delivery date it must be BA compliant. While there was discussions of a waiver, according to the Federal Register, this waiver has yet to be created. Thus local institutions should be BA compliant regardless of when the policy was enacted. [9]

On the supplier side, some firms adjusted in response to the policy, rather than before the policy was enacted. One such firm is New Flyer Industries (NFI), a Winnipeg-based Transit Bus manufacturer[10] and the largest firm according to market share.[11] In a 2018 Mass Transit Magazine article titled *New Flyer Opens High Tech Bus Part Fabrication Facility in Shepherdsville, KY*, the article states "The Shepherdsville facility furthers NFI's effort to insource part fabrication capability, and increases the company's commitment to meet increased US Content requirements under BA provisions of the 2015 FAST Act."(Mass Transit Magazine, 2018). Furthermore, the construction of this facility occurred after an announcement in November 2017, two years after the policy. NFI's factory expansion in Alabama in 2018 is another example of NFI increasing their U.S. presence to comply with content limits (AL.com, 2018).

On the other hand, the railcar industry seems to not have adjusted. A July 2020 article titled "Manufacturing group accuses Metro of sidestepping 'Buy America' clause for new rail cars" by The Washington Post details the debate between transportation agencies and

[9]The lack of new BA waivers since 2016 provides further evidence for the lack of an exemption waiver, see the FTA website here: https://www.transit.dot.gov/regulations-and-guidance/buy-america/waivers-granted

[10]Although NFI is a Canadian based company, much of their business is within the United States. Firing decisions during the Covid-19 pandemic provide further evidence on the importance of the United States Market, see "Busmaker NFI laying off hundreds in Winnipeg due to pandemic", CTV News Winnipeg, March 2020 at https://winnipeg.ctvnews.ca/busmaker-nfi-laying-off-hundreds-in-winnipeg-due-to-pandemic-1.4865667

[11]As of 2018, NFI has 43% market share in the North American Heavy Duty Transit Bus industry. See https://www.nfigroup.com/site-content/uploads/2019/08/NFI-Investor-Presentation-August-22-2019.pdf

manufacturers. While Metro pursued funding opportunities unrelated to BA, the reason cited was due to, "no U.S.-based manufacturers of subway cars." While there are foreign agencies building local factories such as Kawasaki, "Each of the cars is valued at $2 million and the 8000-series is expected to cost as much if not more." These high costs prompted Metro to seek other funding opportunities and foreign agencies for IFBs much to the dismay of domestic manufacturers. These domestic manufacturers argue that domestic manufacturers can build railcars, stating "There are plenty of qualified rail car manufacturers in the United States who can meet or exceed Metro's needs, who can deliver a great product. And those dollars will be supporting American jobs, as well as an investment in the future of transit in the D.C. area." (George, 2020). These claims suggest the cost of using these domestic producers may be too high.

The question then becomes how did the law affect funding decisions at the local level. Figure 3.3 and 3.4 shows two potential avenues of how procurement decisions play a role. The figure shows the estimated spending for rail cars and transit buses between 2009 to 2018. Federal funds used for railcars increases sharply after 2017, when the policy increased from 60% to 65% domestic. Meanwhile, buses appear to have a modest increase before the policy was enacted and later returned to 2016 levels. These spending habits could be explained by a variety of factors[12]. However both charts match the anecdotal evidence. Buses adjusted before the policy was fully implemented while railcars took too long to adjust, leading to major cost increases. In the context of this chapter, rail car purchases could be dominated by international firms with little United State presence, and the investment costs could be too high for firms to change their type. Buses could have had comparatively smaller type switching costs, leading to more competition and smaller procurement costs overall. These bus industries could have adjusted via factory expansions or new factories much quicker,

[12]Frequency of purchasing could be a factor. Railcars are far less frequent purchases compared to buses, as the average durability of buses ranges from 4 to 15 years, while commuter rail is closer to 10 - 20 years (See DOT 2013 and Schneck 2007). Given that the policy had two years before it came into effect, it is unclear why rail car purchases did not happen before the policy came into effect.

before the policy came into effect.

Figure 3.3: FTA Funds Used for Railcar Purchases

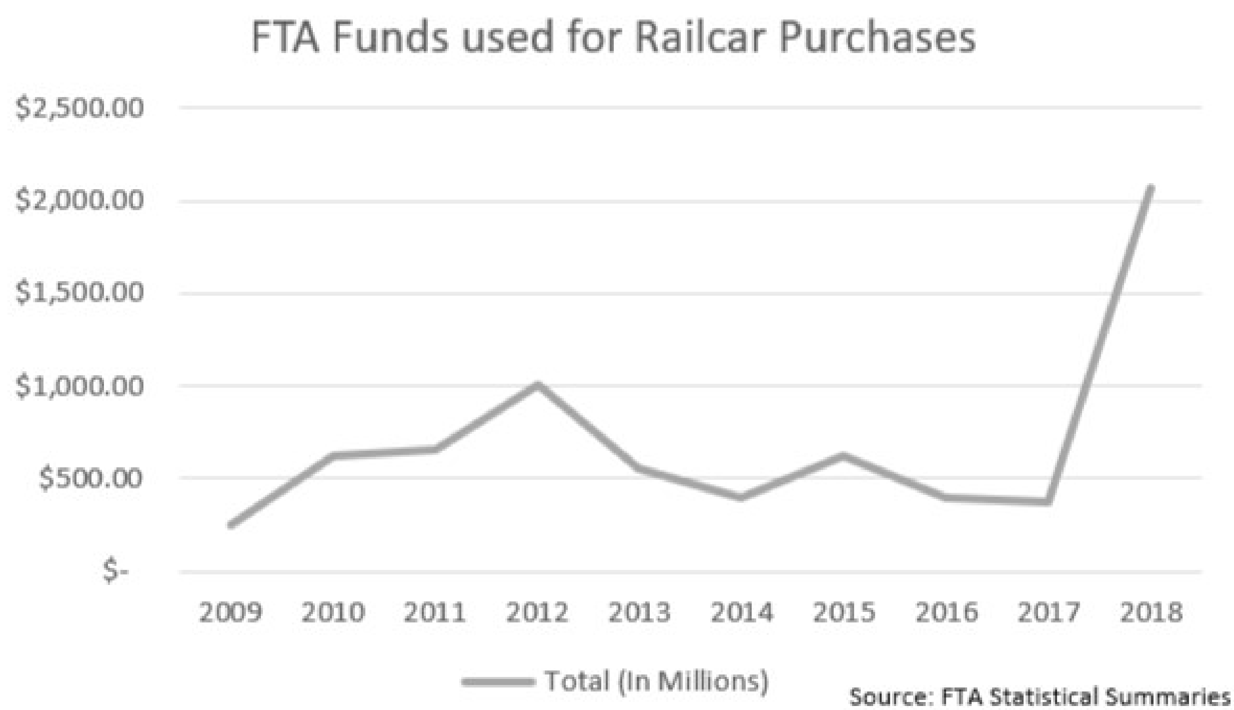

Note: This graph shows the amount of FTA Funds Used for Railcar purchases, collected from the FTA Statistical Summaries. Each year is a Federal Government Fiscal Year, which starts on October 1st and ends on September 30.

This evidence suggests several avenues for which auction micro-level data and a calibrated model can solve. The next section details these models.

3.4 Model

To measure the impact of BA, I use three auction models and consider their implications before and after the policy. I make three central assumptions appropriate for my setting, outlined here. First, auctions for transportation goods are first price, independent values sealed bid auctions. Second, there are two types of bidders, one of which is penalized based on their characteristics. Finally, participation is endogenous and bidders can change their

Figure 3.4: FTA Funds Used for Bus Purchases

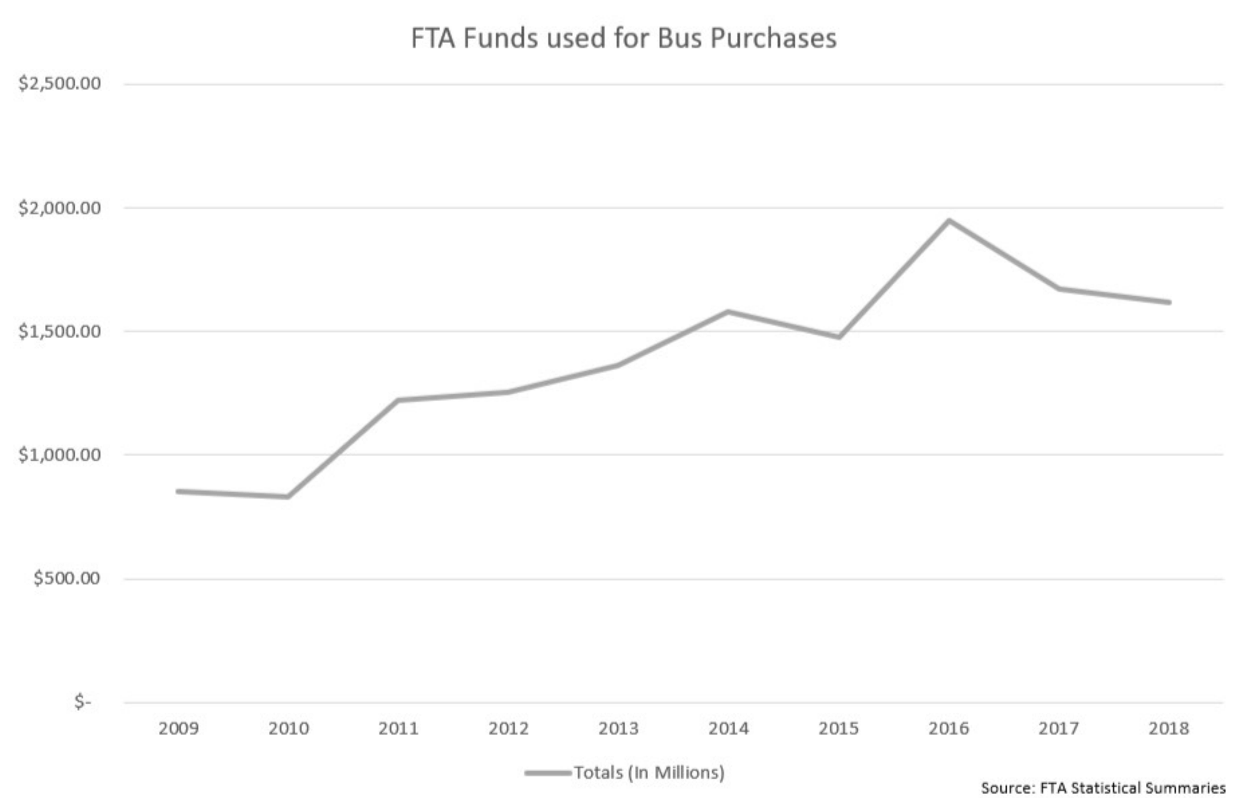

Note: This graph shows the amount of FTA Funds Used for Bus purchases, collected from the FTA Statistical Summaries. Each year is a Federal Government Fiscal Year, which starts on October 1st and ends on September 30.

type. This section outlines the three models that I will later compare.

3.4.1 Model I: Standard First Price Sealed Auction

To start, I consider the standard first price sealed bid auction model without participation effects. The government (the seller), decides to sell an indivisible contract to N risk-neutral firms (the buyers). The contract is awarded to the firm that proposes the lowest price for the contract. Each firm has a cost value c_i drawn from a distribution with cumulative distribution function (CDF) $G(c_i)$ and partial distribution function (PDF) $g(c_i)$. Each firm follows a bidding strategy $b_i = \beta(c_i)$ with bid $b_i \in [\underline{\text{b}}, \bar{b}]$ and has individual cost $c_i \in [\underline{\text{c}}, \bar{c}]$. There is no reservation cost.

The firm's profit is

$$\pi_i(c_i) = (\beta(c_i) - c_i) \times \underbrace{\Pr(c_i \leq c_\ell \forall \ell)}_{\text{Probability of beating all other bidders}} \tag{3.1}$$

Which simplifies to

$$\pi_i(c_i) = (\beta(c_i) - c_i)(1 - G(c_i))^{(n-1)} \tag{3.2}$$

Taking the first order condition with respect to the valuation c_i, I obtain

$$\frac{1}{(\beta(c_i) - c_i)} = \frac{(N-1)(g(\beta(c_i)))}{(1 - G(\beta(c_i)))} \frac{1}{\beta'(c_i) - 1} \tag{3.3}$$

I use the GPV (2000) method of indirectly solving for model parameters. Noting that $b_i = \beta(c_i)$, if $\beta(c_i)$ is invertible, $G(b_i)$ can be mapped to a distribution $F(b_i) = \Pr(c_i \leq \beta^{-1}(b_i)) = G(\beta(c_i))$ with PDF $f(b_i)$. Upon solving the above equation using this distribution, I obtain the solution derived in GPV (2000):

$$\frac{1}{(b_i - c_i)} = \frac{(N-1)(f(b_i))}{(1 - F(b_i))} \tag{3.4}$$

To identify $F(\cdot)$, $f(\cdot)$ and c_i I make the following assumptions. First, $F(b_1, \ldots, b_N) = \Pi_{i=1}^{N} F(b_i)$, or distributions are independent and identically distributed. Second,

$$c_i = b_i - \frac{(1 - F(b_i))}{(N-1)(f(b_i))}$$

is strictly decreasing on $[\underline{b}, \bar{b}]$ and its inverse is differentiable on the support of c_i.

Suppose I have data on the number of firms of each auction, the bids, and the winner's identity. Using this model, I can directly identify the bids and the distribution $(1 - F(b_i))$. The latter can be estimated via the sample analogue

$$(1 - \tilde{F}(b_i) = 1 - \frac{1}{NL} \sum_{l=1}^{L} \sum_{i=1}^{N} \mathbb{1}(b_{il} \leq b)$$

Similarly, the probability density function $f(b_i)$ can be estimated via a kernel density estimator. For convenience, I consider the Epanechnikov estimator

$$K_{ep}(\Psi) = \frac{3}{4}(1 - \Psi^2)\mathbb{1}\{|\Psi| \leq 1\}$$

Where $\Psi = \frac{b_i - B_{pl}}{h_g}$ and h_g is the bandwidth.

In cases where I do not have all data available, I can still estimate parameters within this framework. Consider the case where some bids are unobserved. For example, suppose data is only collected for the top two firms rather than all firms that submitted a proposal. Here, I consider two-player auctions rather than N player auctions. This would lead to weaker estimates of cost distributions, but may be feasible given data constraints.

3.4.2 Model II: FPSA with a Penalized Type

I now consider the framework of Krasnokutskaya and Seim (2011) which considers an auction with heterogeneous types and a penalty on one such type.[13] I include this model to incorporate how this policy effects foreign firms compared to domestic firms.

Firms are split into two types based on their domestic content shares: $k \in \{0\% - 70\% \text{ Domestic}, 70\% - 100 \text{ Domestic}\}$, hereafter denoted as Type 1 and Type 2. Type 1 is affected by the policy, while Type 2 is not.[14] Each type follows a bidding strategy $\beta_k(\cdot)$ with bid $b_i \in [\underline{b}_k, \bar{b}_k]$, has individual cost $c_i \in [\underline{c}, \bar{c}]$, and a mapping from cost to bids given their type designation. The risk-neutral firm minimizes cost, but is subject to federal regulation on where funds can be spent. These regulations are exogenous and induce a penalty on Type 1. The firm wins if they post a bid below all other bids adjusted by the bid penalty δ.

The profit $\pi_i(c_i)$ of the firm is dependent on their probability of winning against their own type and the other type:

$$\pi_i(b_i, c_i) = (b_i - c_i) \times \underbrace{\Pr(b_i \leq b_\ell \forall \ell | k(\ell) = k(i))}_{\text{Probability of competing against own type}} \times \underbrace{\Pr(b_i \leq (1+\delta)^{1-(2\times \mathbb{1}\{k=1\})} b_\ell \forall \ell | k(\ell) \neq k(i))}_{\text{Probability of competing against other type, including bid penalty}} \tag{3.5}$$

Where $\mathbb{1}\{k = 1\}$ is the indicator function for if the firm belongs to Type 1. The above simplifies to

[13] While the full model within Krasnokutskaya and Seim contains endogenous participation and unobserved heterogenity within buyer cost, I save the former for Model III and discuss unobserved heterogenity in a later section.

[14] For the purposes of illustration I choose one type that can potentially be affected by the policy and one type that will not. This framework can be extended to n types as well. I discuss additional types in a later section.

$$(b_i - c_i)(1 - F_C^k(\beta_k^{-1}(b_i))^{(n_k-1)}(1 - F_C^{-k}(\beta_k^{-1}((1+\delta)^{1-2\times\mathbb{1}\{k=1\}}b_i)))^{n-k}$$

Where F_C^{-k} is the cost CDF for type k bidders. The first order condition follows:

$$\begin{aligned} \frac{1}{b_i - c_i} = \underbrace{\frac{(n_{k(i)}-1)f_c^{k(i)}[\beta_{k(i)}^{-1}(b_i)]}{(1-F_C^{k(i)}[\beta_{k(i)}^{-1}(b_i)])}\frac{\partial\beta_{k(i)}^{-1}}{\partial b_i}}_{\text{Optimal Markup versus Own Type}} + \\ \underbrace{\frac{n_{k(-i)}(1+\delta)^{1-2\mathbb{1}\{k=1\}}f_C^{k(-i)}[\beta_{k(-i)}^{-1}((1+\delta)^{1-2\mathbb{1}\{k=1\}}b_i)]}{(1-F_C^{k(-i)}[\beta_{k(-i)}^{-1}((1+\delta)^{1-2\mathbb{1}\{k=1\}}b_i)])}\frac{\partial\beta_{k(-i)}^{-1}}{\partial b_i}}_{\text{Optimal markup versus other type}} \end{aligned} \tag{3.6}$$

This model introduces two major differences. The first is asymmetric bidding behavior. Since I assume foreign firms and domestic firms have different distributions, this changes the identical distribution assumption outlined in Model I. The second difference is that the bids include the bid penalty as part of the optimal function. This leads to situations where Type 1 firms could never win. For example, if there is a Type 1 firm with cost $c_i \in [\frac{\bar{b}_2}{(1+\delta)}, \bar{c})$, if there is a Type 2 firm in the auction then Type 1 firm can never win as the penalty will be too costly for them to overcome.

There exists two boundary conditions that can be used to prove equilibrium existence:

1. Right-boundary condition: If a Type 2 firm has cost $\bar{c}$, they bid $\bar{b}_2 = \bar{c}$ if $n_2 > 1$. If $n_2 = 1$, they have no competition and choose $\bar{b}_1$ such that

$$\pi_i = (\bar{b}_2 - \bar{c})\left(1 - F_1\left(\frac{\bar{b}_2}{(1+\delta)}\right)\right)^{n_2} \tag{3.7}$$

Type 1 firms with $c_i \in [\frac{\bar{b}_2}{(1+\delta)}, \bar{c})$ will never win and bid their cost.

2. Left-boundary condition: There exists a bid level $\underline{b}_2$ such that for all Type 2 firms,

$\beta_2(\underline{c}) = \underline{b}_2$. For Type 1 firms, $\beta_1(\underline{c}) = \underline{b}_1 = \underline{b}_2/(1+\delta)$

Given these conditions Reny and Zamir (2004) establish existence and uniqueness of a bidding equilibrium.

Although firms are asymmetric, Athey and Haile (2002) discuss how the model can be identified and tested in this particular case. If the firm's Type is observed, and assuming $F_C^{k(i)}(\cdot)$ is on the same support for all i, Theorem 2 of Athey and Haile (2002) states the distribution is identified. For the model to be testable, I would need the exogenous variation imposed by the regulatory change and participation differing between auctions. Similar to Model I, Theorem 6 of Athey and Haile (2002) explains that the model can still be testable if there is only data on the top two firms.

3.4.3 Model III: Endogenous Participation and Type Choice in a FPSA with a Penalized Type

I now consider the full version of the model, incorporating the simultaneous entry framework from Roberts and Sweeting (2013) and the foreign type switch. This model carries over Model II's framework: There is a type-specific optimal bidding function, firms know their valuation in the bidding stage, and there is still a penalty on one of the firm types. However, I now impose two preliminary stages before the auction stage: the type-switch stage and the participation stage. The order is as follows:

1. Given a noisy signal of their cost c_i, the foreign firms choose their type - either remain foreign firms (Type 1), or pay some firm-specific cost ξ_i to become domestic firms (Type 2).

2. Given a noisy signal of their cost, and having observed which foreign firms chose to be domestic, domestic and foreign firms choose to enter the auction and pay a fixed cost k.

3. The auction proceeds as follows in Model II.

I include the two stages simultaneously for two reasons: participation effects are vital for understanding public procurement, and the type-switching behavior matches an objective of the policy - making foreign firms increase their domestic capacity. Since Samuelson's (1984) work on auctions with entry costs, participation has increasingly been incorporated into these models. Endogenizing participation is key to understanding the results, especially in government procurement where firms may find the process too costly to participate. Additionally declining participation from foreign firms as the penalty becomes more binding may explain distribution changes. Allowing firms to change their type allows them to circumvent this penalty policy at an ex-ante cost. Whether they invest and how this affects the domestic firms who will have stronger competition is central to the model's implications.

Since the final stage is the same as Model II, I consider the participation stage. Here, firms decide on entry based on what their signal is. The signal is given by $S_{ik} \sim H(\cdot)$, with PDF $h(\cdot)$ and S_{ik} is I.I.D. across bidders. From here, a firm decides what their valuation could be given Bayes Rule. As noted in Roberts and Sweeting, as the variance of $S_{ik} \to 0$, the model resembles Samuelson (1985), where firms know values prior to paying an entry cost. Define the conditional density $f_C^{k(i)}(C|S_i)$ as the conditional density for cost given type $k(i)$ and signal S_{ik}.

In this stage, firms of both types enter if and only if they receive some signal above a cutoff signal S_{2k}^*. At this signal, their expected profit is equal to the entry cost ξ_2:

$$\int_{\underline{\mathrm{C}}}^{\bar{C}} \left[\int_{\underline{\mathrm{C}}}^{\bar{C}} \pi_i(X,Y)dX \right] f_C^{k(i)}(Y|S_{2k}^*)dY = \xi_2$$

Including profit, we have

$$\Rightarrow \int_{\underline{C}}^{\bar{C}} \left[\int_{\underline{C}}^{\bar{C}} (X - Y)(1 - F_C^k(\beta_k^{-1}(X))^{(n_k-1)} \right.$$

$$\left. (1 - F_C^{-k}(\beta_k^{-1}((1+\delta)^{1-2\times\mathbb{1}\{k=1\}}X)))^{n-k} dX \right] f_C^{k(i)}(Y|S_{2k}^*)dY = \xi_2 \tag{3.8}$$

Although the cost is constant across firms, the optimal signal for participation differs across firm type. This is due to both the penalty for one type and the differences in distributions between the two types. For identification purposes, I follow Roberts and Sweeting (2013) and assume $S_{22}^* > S_{21}^*$, or that a foreign firm's signal needs to be lower than a domestic firm's signal. While assumed in Roberts and Sweeting (2013), this occurs mechanically in my setting due to the penalty. Because the penalty occurs on foreign firms only, their signal must be sufficiently low such that any valuation they receive will be lower than the high cost domestic bidders.

I now consider stage 1. Here, the foreign firm decides whether to remain a foreign firm or change their type to a domestic firm. If they change their type, they inherit the domestic distribution for cost and receive a new signal $S_{i,-k}$. I assume their prior signal and cost have no affiliation with this new draw. Their expected profit given that they become a domestic firm is their new expected draw for stage 2 minus their firm-specific adjustment cost:

$$\int_{\underline{S}}^{\bar{S}} \left(\int_{\underline{C}}^{\bar{C}} \left[\int_{\underline{C}}^{\bar{C}} \pi_i(X,Y)dX \right] f_C^2(Y|Z)dY \right) h(\cdot)dZ - \xi_i \tag{3.9}$$

The expected profit given that they stay a foreign firm is

$$\int_{\underline{C}}^{\bar{C}} \left[\int_{\underline{C}}^{\bar{C}} \pi_i(X,Y)dX \right] f_C^1(Y|S_{i1})dY \tag{3.10}$$

Therefore, foreign firms will change their type if the expected profit of becoming domestic is larger than the expected profit of staying foreign, I.E.

$$\int_{\underline{S}}^{\bar{S}} \left(\int_{\underline{C}}^{\bar{C}} \left[\int_{\underline{C}}^{\bar{C}} \pi_i(X,Y)dX \right] f_C^2(Y|Z)dY \right) h(\cdot)dZ - \int_{\underline{C}}^{\bar{C}} \left[\int_{\underline{C}}^{\bar{C}} \pi_i(X,Y)dX \right] f_C^1(Y|S_1)dY \geq \xi_i \tag{3.11}$$

Similar to before, there is a cutoff signal S_1^* such that this holds with equality. At this point,

$$\int_{\underline{S}}^{\bar{S}} \left(\int_{\underline{C}}^{\bar{C}} \left[\int_{\underline{C}}^{\bar{C}} \pi_i(X,Y)dX \right] f_C^2(Y|Z)dY \right) h(\cdot)dZ - \int_{\underline{C}}^{\bar{C}} \left[\int_{\underline{C}}^{\bar{C}} \pi_i(X,Y)dX \right] f_C^1(Y|S_1^*)dY = \xi_i \tag{3.12}$$

Note that foreign firms are the only firm types making this decision.

I have three equations defining three cutoff signals by which I must prove equilibrium existence, identification, and testability. Gentry and Li (2014) provides the framework for all three. Given assumptions for stochastic dominance on $f_C^k(c|S)$, the signal distribution having the same support as the value distribution, and independence on (C, S), an interior equilibrium exists. Given my instrument which plausibly creates exogenous changes in entry behavior, Assumption 4 is satisfied and the model can point-identify distributions and entry costs (Proposition 2, Gentry and Li, 2014). Again, in the presence of a lack of data or a weak instrument, bounds can still be constructed on entry costs as mentioned in the paper.

3.4.4 Extensions and Shortcomings

These models can be extended to include other aspects of the data, such as unobserved auction heterogeneity. This unobserved auction heterogeneity could influence cost estimates. Although here I implicitly assume costs are point estimated by all three models, one could take the approach of Athey, Levin and Seira (2011) and other papers in the literature and parameterize the optimal bidding function as a function of observable covariates. Propositions in both Athey and Haile (2002) and Gentry and Li (2014) show construction on bounds and testability in these frameworks. After parameterizing the bidding function, I could estimate bidder valuation on residuals from bidder function regressions.

Additionally, while I only consider two types, I could include additional types. To study transitory dynamics, a type set could be $k \in \{0\%{-}60\% \text{ Domestic}, 60\%{-}70\% \text{ Domestic}, 70\%{-}100\% \text{ Domestic}\}$. This would require additional notation on Model II and incorporate dynamics in the model and adding additional assumptions for Model III's signals.

Finally, the model requires minimal to no assumptions on distributions. The signal only requires the assumption that it must have the same support as the value distribution. One could test various forms of the signal distribution for robustness. Similarly, one could use a different kernel for estimation for robustness.

Although I do not include domestic to foreign type firm switches for ease of calculation, I can test for evidence of this empirically. Domestic firms could leverage their status and minimize costs through outsourcing, changing their type to foreign until the policy comes into effect. One possibility of testing this would be to use BA Audit data to look for evidence of bunching near the content limit. If domestic firms seem to produce consistently near the penalty thresholds, the model may need to be adjusted to accommodate domestic type switches.

Another fairly strong assumption is the necessity of a new random draw for foreign firms switching types. If a foreign firm decides to switch to a domestic firm classification, they must redraw both a signal and a new valuation. This could be unrealistic, as foreign firms likely have some knowledge about what their shock may be when moving into a new market. While I keep this assumption in for ease of ease of calculation, one robustness check could be to assume foreign firms could take another draw from the foreign distribution but not be subject to the penalty. This would create a third type in the auction that domestic bidders would have to compete against. [15] I refrain from doing that in the baseline model for simplicity, but this could be one potential robustness check.

Data limitations also causes issues. Other papers in the literature with similar models contain at least 600-800 auctions to estimate bidding behavior. The amount of data needed to estimate this model will likely need to exceed this to estimate other parameters of the model, such as firm type switching behavior. Given additional assumptions there may be testable implications, but I present a model with simulated data as a starting framework.

3.5 Estimation and Data

Absent a true empirical framework, I present two model with simulated data to illustrate the ambiguous effects of the policy.[16] I discuss the model outcomes and interpret the results below.

3.5.1 Model I

For the purposes of this model, consider two domestic and two foreign firms competing for a government procurement contract. All firms have a normal distribution $TN(\mu_k, 150)$ truncated at 0. For foreign firms let $\mu_k = 1000$. For domestic firms, let $\mu_k = 1250$, where $k \in \{60\% - 70\% \text{ Domestic}, 70\% - 100\% Domestic\}$.[17]. I first consider pre-policy conditions where Type 1 is not penalized. Figure 3.5 shows the results for domestic and foreign firms.

I now consider the policy. Foreign Bidders now have their bidding costs increased by 0.25%, since they are no longer BA compliant. This will increase the bids. However, in estimation, if the penalty is not considered, the value results will be biased upwards. Figure 3.6 illustrates this.

[15]There may be unobserved heterogeneity between firms that remain foreign versus those that invest and compete as domestic producers, see Fabling and Sanderson (2012)

[16]For plans for the empirical implementation, see the Appendix.

[17]I consider this grouping for ease of exposition.

Figure 3.5: Model I Cost for Domestic Bidders and Foreign Bidders, Pre-Policy

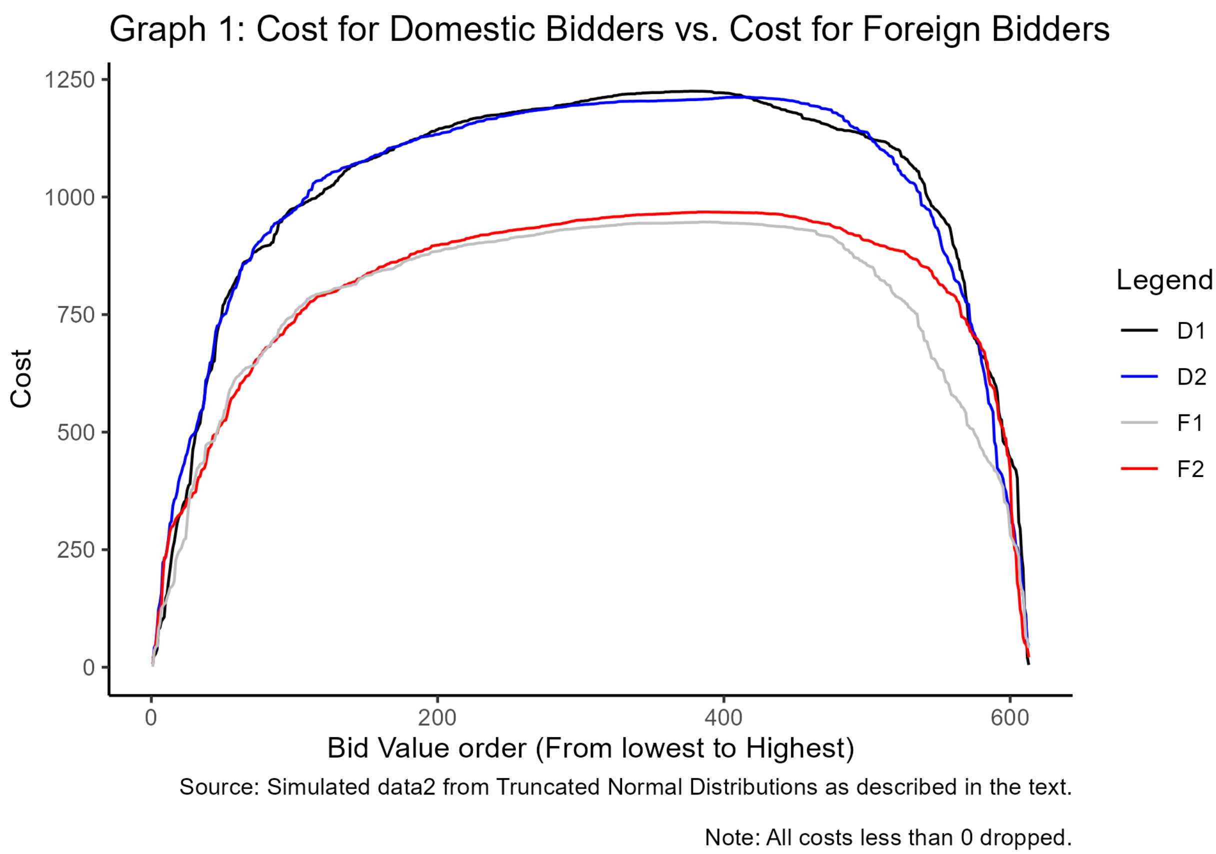

Figure 3.6: Model I Cost for Domestic Bidders and Foreign Bidders, Post-Policy

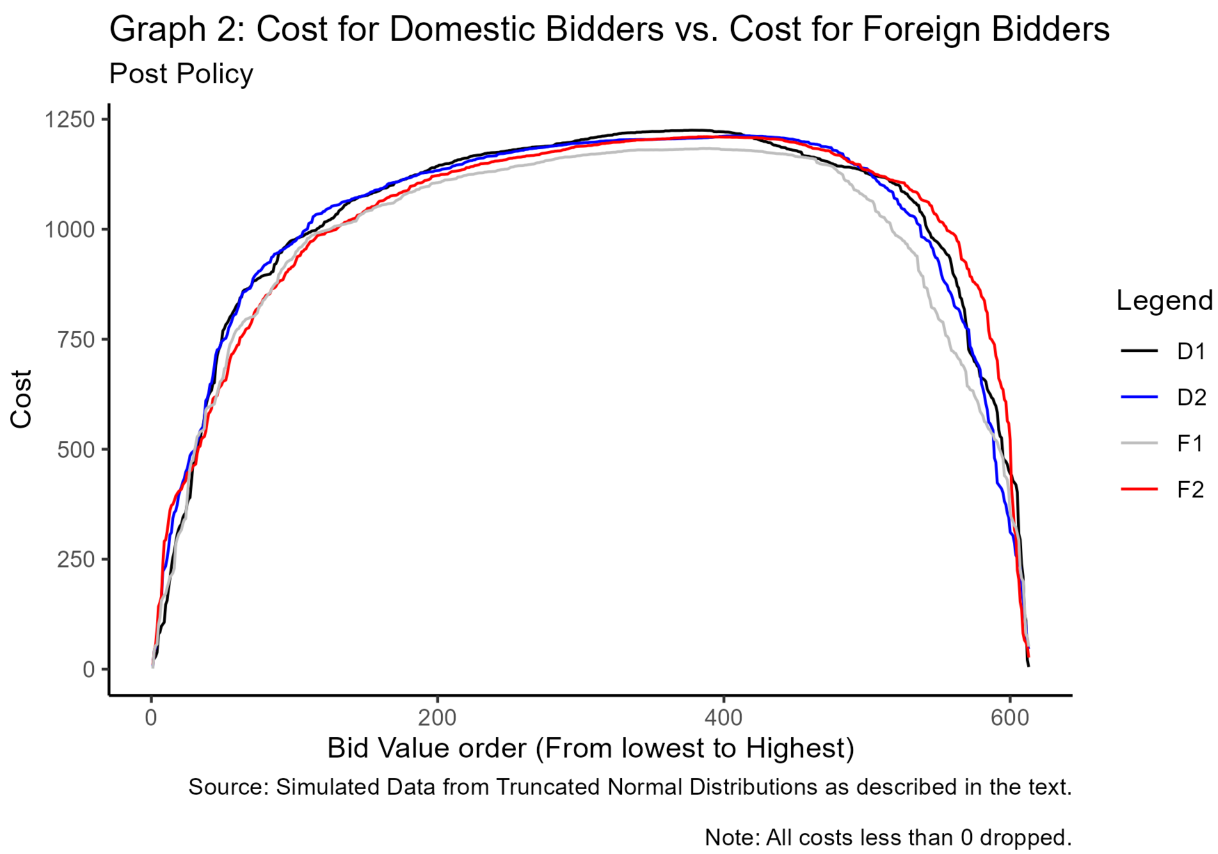

The distribution value has not changed, yet due to the penalty, cost values appear to be equal between types. The foreign cost is biased upwards.

For estimating welfare effects, consider the cost of procurement, the markup for the winning firm and the foreign firm's aggregate probability of winning. Table 3.1 below describes these for the pre- and post- periods.

Table 3.1: Welfare Pre-and Post-Policy Under Model I

	Variable	Pre_Policy	Post_Policy	Difference
1	Mean_Cost	967.65	1182.88	215.23
2	Mean_Markup	200.77	333.23	132.46
3	Percent_Foreign	100.00	56.12	-43.88

Model I's welfare estimates would imply an increased markup due to the policy, an increased cost to the government, and a decline in the percentage of foreign winners. These estimates do not take the penalty associated with being a foreign firm into account.

3.5.2 Model II

I now consider Model II, where firm types are differentiated. I keep the same setting as before. Note that since all firms are not subject to the penalty before the policy, they are effectively the same type in terms of the auctioneer. Therefore, the parameters and the welfare effects for the Pre-Period are unchanged. However, the post period now accounts for the penalty. Instead of having a heterogeneous bidding function, each type now changes their bidding function based on who their competition is.[18] Figure 3.7 displays the post period when the penalty is taken into account:

Figure 3.7: Model II Cost for Domestic Bidders and Foreign Bidders, Post Policy

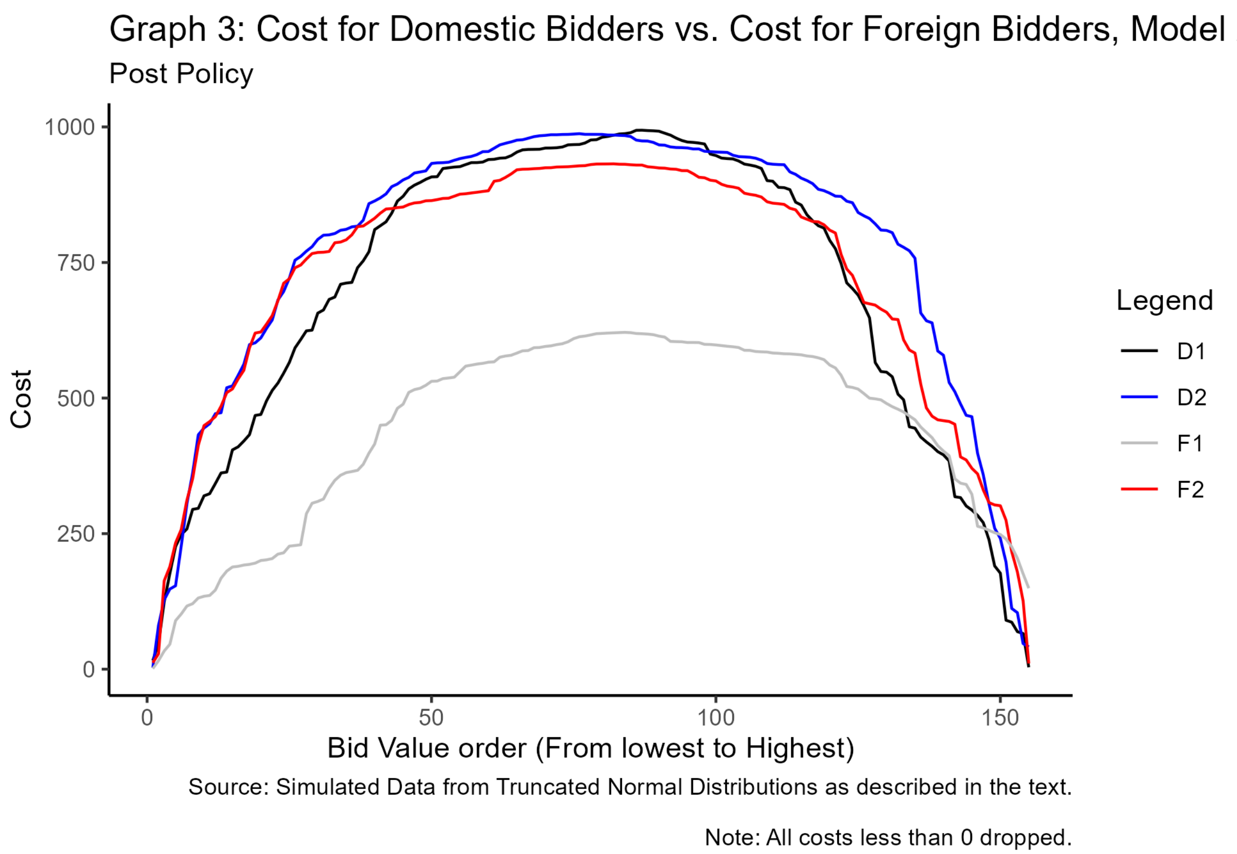

Here, foreign bidder's lower costs are taken into account. For example, F1, one of the foreign firms with the lowest cost, has a more accurate valuation than in Model I. The more interesting results are what a model like this implies for welfare, as shown in Table 3.2:

In this setting, note the mean cost and the percent foreign results are roughly the same

[18]Here, I denote the bidding function as each player stating the bid listed. A natural extension would be to consider different types of bidding functions, but for ease of exposition I consider the simple example here.

Table 3.2: Welfare Pre-and Post-Policy Under Model II

	Variable	Pre_Policy	Post_Policy	Difference
1	Mean_Cost	967.65	1205.95	238.30
2	Mean_Markup	200.77	699.03	498.26
3	Percent_Foreign	100.00	54.19	-45.81

as Model I. However, the markup is much greater in the post-policy period. This is due to the domestic firms incorporating the penalty into their bidding functions. Since domestic firms are not subject to the penalty, they optimally shade their bids upward to gain more profit. Markups increase more compared to Model I.

3.5.3 Model III

I provide an outline of a simulated version of Model III here. Suppose I observe the four firms as before in the Pre-Period. In the Post-Period, I observe that a foreign firm changed their type and a domestic firm dropped from the sample. Due to the lack of endogenous participation in Model II, there will be no explanation for why the domestic firm dropped from the sample, nor why the foreign firm changed their type. Model III deals with each of these issues. To solve the model, one would have to impose a distribution on signal parameters and investment costs. One would then have to use a GMM estimator for a grid search to find cutoff signal parameters such that participation, investment, and bidding behavior fits best.

One proposed distribution for investments and signals is a uniform distribution over the foreign firm's valuation from $[0, \sigma_i V_i]$, where $\sigma_i > 1$. Here firm's investment costs would be a share based on their valuation. Similarly, the participation cost would be the average of these values for both domestic and foreign firms.

While the welfare implications will be similarly calculated as in Model II for the parameters listed, the gain from this model would be to better estimate investment and participation costs. These also explain valuations better, as bidding functions are monotone in the number of bidders. If the number of bidders decreases, the markup shading increases, decreasing the implied value and making markups larger. Examples of these corrections and how they improve welfare estimates are explained in Krasnokutskaya and Seim (2011) and Roberts and Sweeting (2013), among others listed in the literature review.

3.6 Conclusion

This chapter examines possible impact of local content restrictions using the case study of BA's effects on buses and railcars. The paper establishes that the impacts of the content limit increases are ambiguous. Additionally, anecdotal evidence suggests the impacts are largely industry specific and depend on existing domestic plant capacity. This chapter then explores model implications in an exercise involving hypothetical auction data. It finds that Model I compared to Model II understates the markup, but in situations where results may be more ambiguous, Model III can provide guidance and potential explanations for outcomes.

In later versions of this working paper, I plan on providing an empirical foundation to the models instead of using simulated data. This empirical framework will apply these models to transit procurement data and examine outcomes such as the welfare metrics listed above and others implied in Model III, such as investment and participation. Plans for this empirical foundation are located in the Appendix.

3.7 Appendix

3.7.1 Empirical Model

For estimation, I plan on using the techniques proposed in GPV (2000), Krasnokutskaya and Seim (2011), and Roberts and Sweeting (2013) for Model I, II and III respectively. For Model III, I will have to assume distributions for participation, signal, and costs. I plan on using a similar framework for Roberts and Sweeting (2013) to solve for this. This is highly reliant on what data I will have and what initial results from Model I and II resemble.

As for data, I plan on using self-collected data from U.S. transit offices on bus and rail auctions between 2010 - 2024. I additionally plan on using any pre- and post- audit data for BA, if available.[19] Ideally, I would collect data on IFB proposals from each agency that submitted a proposal, their bids, the identity of the bidders, and the identity of the winner of the auction.

One potential concern may be the identity of the bidder - a company, although founded in another country, could still be treated as a domestic firm through initial domestic investments. To resolve this problem, I plan on using the audit data to see how much of the production occurred in the United States. Otherwise, some estimates of how many factories the firm has in the states may be necessary.

[19]There may be other issues with collecting transit authority data. A paper examining the Swedish Bus System noted "While there are some [transportation agencies] that deliver the complete requested documents immediately or within a couple of days, the majority is either slow in handling the request..., deliver incomplete information, claim the documents are lost, or some combination of these." (Vigren, 2017)

3.8 References

Anon. 2020. "New Flyer Opens High Tech Bus Part Fabrication Facility in Shepherdsville, KY." Retrieved July 7, 2020 (https://www.masstransitmag.com/home/press-release/12429289/new-flyer-new-flyer-opens-high-tech-bus-part-fabrication-facility-in-shepherdsville-ky).

Anon. 2018. "New Flyer Opens $25 Million Plant Expansion in Anniston - Al.Com." Alabama News. Retrieved July 8, 2020 (https://www.al.com/business/2018/11/new-flyer-opens-25-million-plant-expansion-in-anniston.html).

Arozamena, Leandro and Estelle Cantillon. 2004. "Investment Incentives in Procurement Auctions." *Review of Economic Studies*, 71(1), 1–18. https://doi.org/10.1111/0034-6527.00273.

Athey, Susan and Philip A. Haile. 2002. "Identification of Standard Auction Models." *Econometrica*. 70(6):2107–40.

Athey, Susan, Jonathan Levin, and Enrique Seira. 2011. "Comparing open and Sealed Bid Auctions: Evidence from Timber Auctions." The Quarterly Journal of Economics, 126(1):207–257. https://doi.org/10.1093/qje/qjq001.

Bajari, Patrick, Stephanie Houghton, Steven Tadelis, Liran Einav, Phil Haile, Igal Hendel, Ali Hortaçsu, Tom Hubbard, Jon Levin, Ken Hendricks, Han Hong, and Harry Paarsch. 2014. "Bidding for Incomplete Contracts: An Empirical Analysis of Adaptation Costs." *American Economic Review*. 104(4):1288–1319.

CBCNY. 2020. "9 Things New Yorkers Should Know About How New York City Buys Stuff." Retrieved August 7, 2020 (https://cbcny.org/research/9-things-new-yorkers-should-know-about-how-new-york-city-buys-stuff).

Fabling, Richard, Arthur Grimes, and Lynda Sanderson. 2012. "Whatever next? Export

Market Choices of New Zealand Firms." *Papers in Regional Science 91(1):137–59*.

Federal Transit Administration, Dot. 2020. FACT SHEET: BUY AMERICA 5323(J).

Federal Register. 2016. "Notice of Proposed Policy Statement on the Implementation of the Phased Increase in Domestic Content Under the Buy America Waiver for Rolling Stock." Federal Register 20049–51.

Gentry, Matthew and Tong Li. 2014. "Identification in Auctions with Selective Entry." *Econometrica* 82.1: 315-344.

George, Justin. 2020. "Metro Accused by Manufacturing Group of Sidestepping 'Buy American' Clause for New Rail Cars." Washington Post. Retrieved August 7, 2020 (https://www.washingtonpost.com/transportation/2020/07/31/manufacturing-group-accuses-metro-sidestepping-buy-america-clause-new-rail-cars/#comments-wrapper).

Guerre, Emmanuel, Isabelle Perrigne, and Quang Vuong. 2000. "Optimal Nonparametric Estimation of First-Price Auctions." *Econometrica*. 68(3):525–74.

Jehiel, Philippe and Laurent Lamy. 1981. "On Discrimination in Auctions with Endogenous Entry." *American Economic Review*. 105(8):2595–2643.

Krasnokutskaya, Elena and Katja Seim. 2011. "Bid Preference Programs and Participation in Highway Procurement Auctions." *American Economic Review*, 101 (6): 2653-86.

Li, Tong and Xiaoyong Zheng. 2009. "Entry and Competition Effects in First-Price Auctions: Theory and Evidence from Procurement Auctions." Review of Economic Studies 76:1397–1429.

Lewis, Gregory and Patrick Bajari. 2011. "Procurement Contracting With Time Incentives: Theory and Evidence." *The Quarterly Journal of Economics*, 126(3), 1173–1211. https://doi.org/10.1093/qje/qjr026

McAfee, RP, J. McMillan-The RAND Journal of Economics, and undefined 1986. n.d.

"Bidding for Contracts: A Principal-Agent Analysis." JSTOR.

Naegelen, Florence, and Michel Mougeot. 1998. "Discriminatory public procurement policy and cost reduction incentives." *Journal of Public Economics* 67.3: 349-367.

Platzer, Michaela D. and William J. Mallett. 2019. "Effects of Buy America on Transportation Infrastructure and U.S. Manufacturing: Policy Options."

Reny, Philip J., Shmuel Zamir CNRS, and France EUREQua. 2002. On the Existence of Pure Strategy Monotone Equilibria in Asymmetric First-Price Auctions.

Miyagiwa, Kaz. 1991. "Oligopoly and Discriminatory Government Procurement Policy." *American Economic Review* 81, no. 5: 1320–28. http://www.jstor.org/stable/2006920.

Roberts, James W., Andrew Sweeting, John Asker, Brendan Daley, Ken Hendricks, Jakub Kastl, Greg Lewis, David Mcadams, Brian Mcmanus, and Rob Porter. 2013. "When Should Sellers Use Auctions?" *American Economic Review* 103(5):1830–61.

Samuelson, W. 1984. "Bargaining under Asymmetric Information." *Econometrica* 52, no. 4: 995–1005. https://doi.org/10.2307/1911195.

Schneck, Donald. 2007. "Useful Life of Transit Buses and Vans". No. FTA-VA-26-7229-07.1.

Federal Transit Administration, D. O. 2017. Conducting Pre-Award and Post-Delivery Audits for Rolling Stock Procurements Best Practices Handbook for Recipients, Auditors, Manufacturers, and Suppliers.

Vigren, Andreas. 2017. "How Many Want to Drive the Bus? Analyzing the Number of Bids for Public Transport Bus Contracts How Many Want to Drive the Bus? Analyzing the Number of Bids for Public Transport Bus Contracts." *Transport Policy*, 72, 138–147. https://doi.org/10.1016/j.tranpol.2018.08.008

White, Richard A. 2016. American Public Transportation Association, Notice of Proposed Policy Statement on the Implementation of the Phased Increase in Domestic Content

Under the Buy America Waiver for Rolling Stock.