



Munich Personal RePEc Archive

Winning Big: Scale and Success in Retail Entrepreneurship

Hollenbeck, Brett and Giroldo, Renato

UCLA Anderson, UCLA

20 May 2020

Online at <https://mpra.ub.uni-muenchen.de/100766/>
MPRA Paper No. 100766, posted 01 Jun 2020 04:51 UTC

Winning Big: Scale and Success in Retail Entrepreneurship

Renato Zaterka Giroldo

UCLA *

Brett Hollenbeck

Anderson School of Management

UCLA[†]

May, 2020 [‡]

Abstract

In this paper we study a novel setting where firms were randomly allocated differently sized retail chains in a new and rapidly growing industry. Beginning in 2014, Washington State used a lottery to allocate licenses to firms in the newly legalized retail cannabis industry. This lottery generates random variation in firm size and in the level of market concentration. We also observe detailed data on all subsequent industry transactions, including prices, wholesale costs, markups, and product assortments. We find that firms that are randomly allocated more retail store licenses in the lottery ultimately earn much higher per store profits than single-store firms. Retailers in multi-store chains charge lower margins, offer larger product assortments, and pay lower wholesale prices. They also face higher but more elastic consumer demand. Similarly at the market level, more concentrated markets have lower average prices and markups. We conclude that higher retail scale and a more concentrated retail sector can benefit consumers and firms alike.

Keywords: Economies of scale, retail pricing, markups, entrepreneurship

*rgiroldo@ucla.edu

[†]brett.hollenbeck@anderson.ucla.edu

[‡]The authors wish to thank the Morrison Center for Marketing Analytics for generous funding. We also wish to thank seminar participants at SMU Cox School of Business, Washington University in St. Louis Olin Business School, EARIE, and UCLA Anderson School of Management for helpful comments.

This paper studies a novel setting where entrepreneurs were randomly allocated differently sized retail chains in a new and rapidly growing industry. We study the effects of a retail chain’s size on both its profits and the price and assortment strategies it employs, as well as the effects of market level concentration on prices and markups.

Firm size is generally considered a key determinant of both firm choices and outcomes in industrial organization, trade, macroeconomics, and finance.¹ Despite this, there is very little evidence on direct causal effects related to firm size because a firm’s size is just one outcome determined in equilibrium along with many endogenous choices made by firms. A retailer’s size or scale both cause and are determined by its pricing policies, its assortment choices, its managerial quality, its capitalization, its set of upstream relationships, its degree of competition with rivals, and so on. Scale is often taken as an explanatory variable for these types of outcomes despite them all being determined in a coordinated fashion. We take advantage of a novel natural experiment that generates direct exogenous variation in firm size to overcome these.

This natural experiment also generates direct exogenous variation in the level of market concentration. There is growing evidence that economy-wide markups have increased over the past few decades and many have speculated about the role of growing firm size and market concentration in driving this increase (De Loecker and Eeckhout (2018) and De Loecker et al. (2020), Grullon et al. (2019), Autor et al. (2020)). Nevertheless, there have been few industry specific studies addressing this issue (Syverson (2019)). This debate suffers from the lack of direct causal evidence on the relationships between firm size, concentration and markups.

The setting we study is the newly legalized retail cannabis industry in Washington state, which began in 2014 and features a number of advantages as a laboratory for the study of retail, entrepreneurship, firm size, and market concentration. First, the number of new firms allowed to enter was capped by regulatory design and excess demand for entry licenses by entrepreneurs led to a lottery to allocate them. Firms could win multiple licenses in this lottery such that two firms

¹Choices and outcomes include pricing as we discuss at length but also investment, innovation, tendency to export, product variety offered, and many others. While much of this research on economies of scale focus on manufacturing industries, economies of scale in retail have also drawn a great deal of attention, including: Foster et al. (2006), Foster et al. (2016), Hortacsu and Syverson (2007), Bronnenberg and Ellickson (2015), and Ratchford (2016) who discuss a number of issues and trends around growing scale in retail. Holmes (2011) and Ellickson et al. (2013) estimate the cost-side benefits of scale for retail chains. Hosken and Tenn (2016) provides an overview of the issues in studying merger and consolidation effects in retail. Rhodes (2015), Armstrong et al. (2009), Rhodes and Zhou (2019), and Moraga-Gonzalez and Petrikaite (2013) provide theoretical studies of consumer search and multiproduct firm strategy. These studies consider retail consolidation and find ambiguous predictions on prices and consumer welfare from larger retailers or retail mergers.

applying for the same number of licenses would, purely by random draw, end up with different numbers of stores and thus different post-entry scale. This lottery that allocated licenses to own and operate retail stores is essentially unprecedented and offers a unique opportunity to study the role of retail chain scale on outcomes in a transparent way.

Second, because the market is closely monitored by regulators, there exists exceptionally good data on post-entry outcomes. We observe every transaction conducted in the industry starting with the first sales, including upstream transactions. This means we directly observe retail prices, store product assortments, vertical arrangements between retailers and manufacturers, wholesale prices, and markups, all at the transaction level. These stores do substantial amounts of sales, averaging \$2.2 million per year in revenue. Third, the industry is new and therefore we observe all entry and the full evolution of firm outcomes over time. This allows us to evaluate whether differences between firms are short-term effects on new entrants or if they persist. The novelty of the market as a whole means retailers are especially important. Consumers must discover what products they value, producers must decide what products to make and how, and retailers act as the intermediaries between these two groups as the industry evolves, deciding what products to stock, what prices to charge, what manufacturers to purchase from, and how to compete with rivals.

Using the retail store lottery for identification, we find that stores that are part of multi-store chains are substantially more profitable than stores operating alone. Their profits are higher by an average of \$380,000 per store per year in the last year of the data, a more than 25% increase. If we simply compared these retailer profits in cross-sectional data, we would see this positive correlation between number of outlets and profits, and therefore we might conclude that higher quality firms earn higher profits and their higher quality also allows them to grow and open more outlets. While this effect is generally likely to be true, our first contribution is to show that there is a direct causal effect as well. Firm size matters, and higher scale causes higher profits. Moreover, for the type of firm operating a mom-and-pop style retail outlet, this difference in profits is enormous.

Our second contribution is to study the mechanism causing larger firms to earn higher profits. Doing so helps shed light on several ongoing debates. By showing direct causal effects of firm size and market concentration on prices and markups we contribute to the growing literature studying the long-term growth of markups and its causes. This debate has primarily been informed by macroeconomic trends, but economy-wide data generally lack the richness to study specific mechanisms and find robust causal effects. Further, models of imperfect competition used in the

macroeconomic markup literature sometimes produce an inherently positive relationship between firm size and markups as well as a positive relationship between market concentration and markups (see Kimball (1995) for example or Mrázová and Neary (2017) for an overview). This literature has also primarily used production function estimates to infer markups, where we can observe them directly.

A view from IO has noted that concentration and markups can increase at the same time under richer models of imperfect competition (Syverson (2019)) and argued for a more nuanced view of the mechanisms generating markups and their implications for welfare (Berry et al. (2019)). At the same time, the IO literature has produced a rich but somewhat contradictory set of predictions on how firm size will effect prices and markups. In particular, the theoretical study of consumer search and retail competition has made great progress in advancing our understanding of the role of the retail sector in markets with consumer search costs, but fundamental issues still remain unresolved. One of these is why different retail stores charge different prices for the same product, and in particular whether larger retailers will ultimately charge higher or lower prices, holding costs and other factors fixed. The disagreement stems from different ways of modeling the nature of consumer search and demand and how these effect retailer pricing incentives.

One view, shown in models such as McAfee (1994) and Armstrong and Vickers (2020) is that larger firms should have more “captive” customers who do not price search and therefore they should charge higher prices than smaller firms.² Another view emphasizes that retailers with more outlets can more easily offer larger assortments, which Anderson and De Palma (2006) demonstrate should result in higher prices.

On the other hand, if the firm with multiple outlets is more prominent but its customers are not “captive”, it should have an incentive to charge lower prices to prevent them from searching at rivals (Armstrong et al. (2009)). Retail stores that offer more products might also attract a more price-sensitive “mass market” set of consumers while small firms are left to offer a more niche assortment but sell at higher prices to high value consumers (Rhodes (2015)). This would lead chain retailers to charge lower prices but earn higher profits. Thus, it is ultimately an empirical question why different stores can charge different prices for the same good, whether large or small

²Prior work has found that both hotel chains and chain restaurants benefit from reputation effects that allow them to charge significantly higher overall prices in this way, see Hollenbeck (2017) and Klopach (2018). Hollenbeck (2017) finds that chain hotels earn roughly 25% higher revenues than otherwise identical independent hotels and attributes this to greater consumer information on chain hotel quality due to their ability to build reputation over multiple outlets. Klopach (2018) finds a similar demand side advantage in the restaurant industry.

firms will charge higher prices, and whether increases in market concentration will cause higher or lower prices.

We therefore investigate the mechanism causing firms with more retail stores to have higher profits and emphasize retail pricing decisions. We find that larger firms charge substantially lower retail prices than smaller firms, including for the same products. Because we observe wholesale prices, we can rule out that this is merely the result of cost-side economies of scale. While we do find evidence of traditional cost-side retail economies of scale, in that the larger firms pay lower wholesale prices for the same products as smaller firms, the difference in retail prices is substantially larger. Larger retailers charge both lower prices and lower margins, in other words. This is despite having larger assortments and presumably enjoying higher awareness or reputation, as well as higher market power.

Among the models predicting lower prices for larger firms, we find that the evidence is more consistent with the view that larger assortments attract more price-sensitive customers (Rhodes (2015)) as opposed to them pricing lower due to a prominence effect. The gap in prices between large and small firms is not initially present but grows substantially over time, as does the difference in assortment size between stores in multi-store chains and stores operating alone.³ The difference in profits between large and small firms is similarly growing over time. Next, we estimate a simple model of consumer demand and find that multi-store firms face significantly more price sensitive customers than do single-store firms and that this difference is growing over time. The evidence suggests that consumer demand shifts outwards for larger firms, as evidenced by their higher profits and sales, but the marginal consumer visiting a multi-store firm is substantially more price sensitive than the marginal consumer faced by a single-store firm. These results together suggest that customer preferences over retailers in this industry are driven more by preference for larger assortments and that this assortment effect has a downward effect on prices (as in Rhodes (2015)) rather than upward (as in Anderson and De Palma (2006)).

Additionally, we test the results at the market level, and measure the causal effect of market concentration on markups and prices. To do so we estimate a simple matching model comparing markets with multi-store chains to those without. Markets are matched using the number of licenses allocated and the number of applications filed, both pre-lottery outcomes, and we argue that,

³A similar type of result is seen in Ilanes and Moshary (2019), who study the deregulation of Washington's retail liquor industry and find that increases in competition lead firms to offer larger assortments.

conditional on these two variables assignment of the treatment is effectively random. Our results show that markets that randomly receive higher levels of concentration have lower average prices and markups than less concentrated markets.⁴

These results have significant implications to the study of concentration in retail markets and about market concentration more broadly. In principle, consumers could either benefit or be harmed by a more concentrated retail sector with larger chain retailers. We find the evidence is more consistent with the view that concentration alone or firm size alone are poor measures of market power. We show that larger firms charge lower prices, lower margins, and offer larger assortments, each of which clearly benefit consumers, and these benefits do not just result from lowering retailer costs. This fact is directly relevant to growing debates over horizontal mergers among competing retailers.⁵ More broadly, our results show that when many strategic variables are available to firms and are chosen simultaneously, and when consumers respond endogenously, large firms can increase their profits substantially in ways other than charging higher markups than small firms. Similarly, more concentrated markets may have lower average prices and markups than less concentrated markets depending on the endogenous choices of larger firms.

While we acknowledge the limitation of our study to relatively small firms in a single industry, industry studies have the potential to inform this debate in ways that aggregate studies cannot. For the broader debates on trends in firm size and concentration and their effects on prices, markups, and welfare, the value of industry studies are that they allow the researcher to model and account for industry idiosyncrasies that aggregate data miss out on. Comparing levels of concentration and prices across industries might miss out on differences in demand and richer strategic responses by firms that are endogenous and interact with pricing. In addition, and most importantly, econometric techniques can be used to help analyze the causal effect between the two variables. And while our

⁴Prior work has shown that markups can fall at the same time that concentration increases. For instance, Syverson (2004a), Syverson (2004b) show that increases in the extent to which consumers can substitute between producers can shift market share to larger but lower cost producers. More closely related is Goldmanis et al. (2010) who show that reductions in search costs brought about by e-commerce can increase concentration and decrease margins in retail settings. Our results and our emphasis on consumer search as a mechanism are consistent with these prior studies, although our empirical design is not to study the effects of changes in search costs but to study changes in size and concentration directly.

⁵A key issue in the analysis of downstream or retail mergers is the concept of countervailing buyer power, that is the idea that an increase in market power downstream might benefit consumers by increasing retailer buyer power relative to suppliers, decreasing input prices, and passing a portion of the savings along to consumers. While we do not literally study a merger, the random variation in number of stores in a chain allows us to analyze the effects of concentration on input and final prices. It is a theoretical debate in what conditions a downstream merger can improve social surplus (Inderst and Shaffer (2007) , Symeonidis (2010), Loertscher and Marx (2019)) and this has limited empirical evidence (Barrette et al. (2020)). We find that greater downstream market power does reduce wholesale costs and by lowering prices and increasing quantity sold also benefits consumers.

study focuses on relatively small retailers, outcomes for small firms and entrepreneurs are often seen as especially important. Our results show that scale has a substantial effect on the success of entrepreneurs. This suggests that barriers to scale, such as capital constraints and legal red-tape have potentially very large costs on entrepreneurs and consumers alike.

The rest of the paper proceeds as follows: Section 1 describes the data and setting with an emphasis on the transaction-level retail data and the retail license lottery, Section 2 provides results on the differences in variable profits and price and assortment decisions between multi-store chains and single-store firms, Section 3 studies consumer demand differences between these types of stores, Section 4 shows results at the market level, and Section 5 discusses and concludes.

1 Data and Setting

This section describes the institutional setting and key features of the data. The regulatory setting dates back to a November 2012 popular referendum passed by voters in Washington state. The new law made marijuana products legal for licensed firms to produce and sell and legal to purchase by any person over 21 years of age. The state legislature subsequently created a tax and regulatory regime for the new legal market by passing I-502 which set up the rules for the legal market to begin sales in July 2014. The state created 3 new types of firm licenses, differentiated by their position in the vertical structure of the industry, similar to the three tier system for alcohol regulation.

Firms can be licensed as retailers, processors, or producers. Processors and producers are allowed to hold both licenses and vertically integrate, but retailers are not allowed to vertically integrate. In addition, the total number of retail licenses was strictly capped. This license cap and how licenses were allocated forms the basis of our empirical strategy and we therefore discuss it at length.

License Lottery:

During the creation of the legal marijuana industry, Washington decided to strictly limit the total amount of entry by retailers. This was motivated by concerns about widespread use of marijuana, which is thought to have negative health effects and social externalities.⁶ In addition, there was a concern about the impact of over-entry by retailers on neighborhood character and property values.⁷ Finally, one of the goals of legalization is to remove marijuana sales from the black market so that

⁶For a review of these issues, see Hall et al. (2019).

⁷See Tyndall (2019) for a study of the effect of dispensaries on nearby home prices, which finds close to zero but potentially small negative effects in Vancouver, BC.

they can be regulated and monitored, and this goal is more easily achieved with fewer retail shops for regulators to monitor.

The result was the choice to limit entry to an initial total of 334 retailers for the state of Washington.⁸ These licenses were allocated at the city level, with the allocation determined by population, population density, and an estimate of past-month marijuana users taken from historical survey data. Expectations by market participants were that this industry would be highly lucrative and demand to enter the industry by entrepreneurs significantly exceeded the number of entrants preferred by the state government, leading to the unusual choice to allocate licenses via a lottery.

Lotteries were held separately at the city level, and 75 cities experienced excess demand for the available retail licenses, resulting in 75 different lotteries being held. In addition, in 48 cities there was not excess demand for licenses. We observe the full list of applicants as well as the ordering determined by the random draw that constituted each lottery.⁹ In order to potentially win a retail license in the lottery, firms needed to file a valid application, which included securing a location for the retail store within the regulatory guidelines and paying a \$250 non-refundable application fee.¹⁰ Among the regulatory guidelines the proposed location of each store had to be at least one thousand feet from elementary or secondary schools, public parks, libraries, among other locations. These regulations along with reluctance by landlords to permit cannabis shops made securing the store location a major barrier to entry for filing an application for the lottery.

Table 1 shows summary statistics on license applications in lottery markets. On average, there were 4.1 applicants per license in these markets with a wide degree of variation. In the largest market, Seattle, there were 191 applications for 21 licenses. In order to prevent large firms from dominating the retail sector, firms were not allowed to own more than 3 retail licenses.¹¹

In July 2015, the state updated its licensing regime. In the first year of the market, Washington had failed to close down formerly unregulated medical marijuana retailers and they were operating in a grey market. In July 2015, Washington increased the total number of retail licenses available from 334 to 556 in order for some of these medical marijuana retailers to enter the market. At the

⁸This number was chosen somewhat arbitrarily to match the number of state-owned liquor stores under Washington's state monopoly on retail alcohol sales that lasted until 2012.

⁹The lottery results for the market of Longview, WA were not available.

¹⁰If issued, the store was responsible for paying a \$1000 annual fee for issuance and renewal.

¹¹According to the regulation, "Any entity and/or principals within any entity are limited to no more than three retail marijuana licenses with no multiple location licensee allowed more than thirty-three percent of the allowed licenses in any county or city." Other regulations included: the prohibition of internet sales and delivery of product, the prohibition of sales across state lines and the sale of marijuana products below their acquisition cost.

same time, any medical marijuana retailer that did not receive a license was forced to close. The newly available licenses were awarded according to the initial lottery draws and thus these draws were still overwhelmingly the mechanism by which the new licenses were allocated in the second wave.

Table 1 also shows the resulting distribution of stores across markets. The number of stores per thousand people in 2017 was 1.38 but with substantial variation. Table 2 shows the joint distribution of applications filed in the lottery and store licenses won. Conditional on the number of applications filed, the number of licenses won is the result of a random draw. This represents the basis of our identification strategy.

There are two complicating factors. First, licenses can be acquired in both non-lottery markets or through the secondary market where firms that own licenses can be purchased by other firms. The second of these is rare and we observe it when it occurs.¹² Second, we observe loose partnerships between stores that are owned by different firms. This typically takes the form of stores operating under similar names or sharing the same website. We carefully document these practices in our data.

Ultimately our goal is to compare the outcomes of stores in multi-store chains to those operating in single-store firms using the lottery to provide cleanly exogenous variation, so multi-store chains that acquire their licenses in these ways or form partnerships would not make for valid comparisons. Therefore we take a conservative approach and exclude any firm who acquired its stores outside the lottery or other multi-store partnerships from our definition of “multi-store”. The result is that some of the firms we designate as “single-store” may actually benefit from the same economies of scale as the multi-store firms, either on the cost side or on the demand side. Therefore all our results should be taken as lower bounds for the true size of these effects. Finally, we note that other than their chain size, the two types of stores are broadly similar on observable characteristics, with no significant differences in their entry timing or location, and single-store and multi-store firms have 13% and 15% of their locations in Seattle, respectively.

We also collect data from Kantar Analytics on advertising spending for firms in this industry. Summary statistics for this data are shown in the Appendix, as well as robustness tests for our main results when advertising spending is included as a covariate. We find little difference in advertising spending between single-store and multi-store firms and little change in our results when advertising

¹²Only 32 licenses are bought or sold during the sample period.

is included.

Transactions Data:

In addition to the awarding of entry rights via lottery, there is a second unique feature of this setting. In order to tightly monitor the marijuana industry, Washington requires all industry participants to enter all transactions into an administrative database. Thus we are able to observe all sales that have ever taken place in the industry, including both retail sales to consumers and sales between retailers and wholesalers, both at the transaction level. These data include both retail price and wholesale price for each transaction. Observing wholesale costs is particularly unusual as this is typically carefully guarded information.

Most notably, we are able to construct profit margins at the transaction level by measuring the wholesale price each retailer paid for each product and link those prices to the final retail sale. In addition, the sum of these margins over all transactions gives a direct measure of variable profits. In total, we observe roughly 80 million transactions worth \$2.5 billion between July 2014 and September 2017. Table 3 presents summary information on the distribution of monthly revenues and variable profits across firms and time.

Retailers in this industry earn large revenues, with a mean of \$180,000 per month or \$2.2 million per year in 2017. This is somewhat right-skewed, as the largest retailers average \$6-10 million per year. To provide a comparison, the average store-level revenue is approximately 3-4 times larger than the average revenue of Washington's liquor stores from 2012-2015, according to Ilanes and Moshary (2019).

Table 3 also shows the distribution of average prices and total monthly sales across firms and time. Figure 1 shows visually how average wholesale and retail prices started out much higher in the first year of the industry, were highly volatile for the first 12 months, and eventually settled down to a stable lower price. In 2017, retailers charged an average price of \$15.2 per unit for marijuana products and paid an average price of \$7.5 per unit to their upstream supplier. This leaves an average profit margin of .54, or a markup of roughly 140%. Because prices are volatile and falling in the first two years of the data as most firms (manufacturers and retailers) are still entering the market, we focus most of our analysis on the last year of the data when the market is relatively stable and mature.

Table 1: Summary Statistics for 74 markets

Variable	Mean	Std. Dev.	Min	Max
Initial Cap	3.43	3.72	1	21
Revised Cap	5.82	6.92	1	42
Applicants per Market	15.07	25.62	2	191
Distinct companies per Market	12.52	18.69	2	135
Applicants per License	4.07	2.78	1.25	16
Prob. Win	0.34	0.17	0.06	0.8
Stores per 1k People ^[1]	0.35	1.65	0.01	29.41

Note: [1] Population count is taken from the 2010 Census. Number of stores from Jan/2017.

Table 2: Joint Distribution of Applications and Licenses

		<i>Approved</i>		
		1	2	3
Applied	1	117	0	0
	2	60	10	0
	3	76	18	3
	4	26	8	0
	5+	22	10	3

Note: This table shows the joint distribution of applications filed and stores ultimately won in the 2014 retail lottery.

2 Effects of Scale on Profits and Prices


The first question we study is whether there exists a direct causal effect of firm scale on profits in the retail sector and how large is it. In general, scale is a strategic choice. We therefore expect

Table 3: Monthly Variable Profits, Revenue, and Total Sales

Subsample:	Mean	Std. Dev.	5th %ile	95th %ile	Mean Seattle	Mean Lottery	Mean Non-Lottery
Variable profit:							
2014	45,508	66,328.6	555.7	179,369.4	112251.7	49,870.6	29,966.2
2015	81,505.8	88,407.5	3,734.2	239,839.5	128,837.8	91,067.3	52,886.7
2016	96,069.7	95,846.3	5,670.3	274,647	109,313.7	116,866	66,536
2017	83,793.1	82,037.6	6,184.9	240,325.5	97,674.2	104,280.7	63,831.4
Revenue:							
2014	81,441	106,468.4	3,274.8	312,072.5	210,699	89,991.1	50,981.5
2015	146,057.4	152,181.1	10,131.5	428,832	233,379.1	162,932.6	95,547.5
2016	184,124.4	180,933.7	13,512.1	540,096.4	204,465.8	222,994.6	128,923.6
2017	170,765	162,839.2	14,538.1	487,896.3	193,003.8	211,094.9	131,470.3
Retail Price:							
2014	17.9	6.4	8.1	29.6	18.9	18.4	16
2015	10.9	2.7	7.3	15.2	11.7	11.1	10.4
2016	8.8	1.7	6.6	11	9.5	8.9	8.8
2017	8.3	1.3	5.8	10.3	8.8	8.3	8.3
Total Units Sold:							
2014	4,541	5,468.9	232	15,213.8	12,529.6	5,106.5	2,526.7
2015	13,801.1	14,590.9	893.4	40,924	20,973.8	15,325.5	9,238.6
2016	21,214.5	21,443	1,564.4	59,753	21,875.8	25,547.5	15,060.9
2017	21,029.9	20,838.4	1,618.7	59,033.6	21,969	26,012.5	16,175.3

Note: This table shows summary data on store-level variable profits, revenues, retail prices, and units sold. Data shown are monthly and do not include the first partial month in which each store opened. Retail price data and total sales are calculated within store each month. Lottery markets refer to markets where excess demand for retail licenses resulted in entry lotteries.

Figure 1: Average Retail and Wholesale Price By Category Over Time (\$/gram)



Note: These figures show category level average retail prices and wholesale prices over time. Retail prices are shown on the left and are tax-inclusive. Wholesale prices are shown on the right.

a positive association between firm size and profits, but this is due in part because more efficient firms are able to grow larger by expanding and adding more outlets. In this case, it is a challenge to separately identify the effect of scale from firm quality in profitability. Our setting removes this aspect of reverse causation and allows us to directly measure the effect of randomly generated size differences on profitability.

An important component of the empirical strategy is that we observe the applications filed that did not win licenses via the lottery. Firms who apply for more entry licenses may have higher quality management, greater commitment to the industry, or be better capitalized than firms applying for fewer licenses. This generates potential concerns that scale itself is not random, but a direct result of the number of applications submitted by the firm. We are able to control for number of applications filed and rely on the retail license lottery to generate random variation in firm size.

2.1 Variable Profits

Our primary dependent variable is the variable profits constructed from the transaction-level sales data. These are defined as total monthly revenues minus total monthly wholesale costs incurred. This differs from profits after other costs such as rents, wages, marketing and other costs. We discuss below how these costs may effect the interpretation of our results but take observed variable profits as the initial benchmark outcome of interest.

Table 4: Effect of Multi-Store Firm Membership on Store Profits

	(1)	(2)	(3)	(4)	(5)
Multi-store (lottery)	29029.4*** (6350.790)	24833.5*** (6446.837)	24456.7*** (6436.918)	32031.1*** (6025.207)	31799.4*** (5512.789)
# Applications FE		Yes	Yes	Yes	Yes
Month-Year FE			Yes		Yes
Market FE				Yes	Yes
Age in Months FE					Yes
Observations	3332	3332	3332	3262	3262
R^2	0.006	0.022	0.029	0.483	0.579

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: This table shows the effects of being in a multi-store firm on store-level variable profits. Only multi-store firms who acquired each of their licenses via the retail lottery are included. Data sample period is April 2016 to April 2017. Standard errors are robust and clustered at the store level.

We ultimately want to identify the causal effect of firm size as measured by whether or not a store is part of a multi-store chain on store-level profits. To measure this effect we regress store profits on an indicator for whether or not the firm won multiple stores in the lottery. Table 4 shows the results of this regression with different fixed effects included. In each case we focus on the last year of the data in recognition of the fact that this is an evolving market with entry of new firms (both at the retail level and upstream) taking place throughout 2014-2016 and prices falling rapidly in 2014-2015. By the last year of the data the industry is more mature and the number of firms and their prices and sales levels are relative stable.

First, in column 1 we show the baseline result with no control variables. It shows that stores in multi-store chains earn roughly \$29,000 more in variable profits than single-store firms. Column 2 includes fixed effects for number of applications filed in order to account for differences in management quality or capitalization that may lead to more applications and thus more stores. The coefficient on the multi-store firm dummy therefore isolates the effect of larger firm size on profits which, conditional on the number of applications filed, is generated by random chance via the lottery. This reduces the profit difference to just under \$25,000 per month. Columns 3 and 4 show the effects of date and market fixed effects, respectively. These fixed effects control for differences across markets in profitability and time trends. Finally, column 5 shows the results with each of these as well as age fixed effects to account for differences in entry timing and the natural increase in profits during a firm's first year. Accounting for each of these results in shows that

Table 5: Effect of Multi-Store Firm Membership on Store Profits

	(1) 2016-17	(2) 2014-2017	(3) 2016-17 Stores \leq 1 year old	(4) 2016-17 Stores $>$ 1 year old	(5) 2016-17 ln(profits)
Multi-store (lottery)	31799.4*** (5512.789)	23706.6*** (4761.062)	24710.7*** (5991.844)	66977.3*** (8972.588)	0.55*** (0.064)
# Applications FE	Yes	Yes	Yes	Yes	Yes
Month-Year FE	Yes	Yes	Yes	Yes	Yes
Market FE	Yes	Yes	Yes	Yes	Yes
Age in Months FE	Yes	Yes	Yes	Yes	Yes
Observations	3262	5523	1530	1817	3262
R^2	0.579	0.513	0.596	0.633	0.540

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: This table shows the effects of being in a multi-store firm on store-level variable profits. The dependent variable is variable profits in columns 1-4 and log of variable profits in column 5. Data sample period is April 2016 to April 2017 in columns 2-5 and 2014-2017 in column 1. Columns 3 and 4 show results only for stores less than one year old greater than 1 year old, respectively. Standard errors are robust and clustered at the store level.


after conditioning on number of applications as well as market, time, and age, stores in multi-store chains earn roughly \$32,000 higher monthly profits than single-store firms. The effect is highly significant and its size amounts to roughly \$380,000 in higher annual profits. This difference is quite substantial, as median store profits during this period are \$820,000 per year and mean store profits are just under \$1.2 million.

Next, we take our preferred specification with all fixed effects included and show how the effect varies over different subsamples of the data. Column 1 repeats the main result from Table 4 that used only the last year of the data. Column 2 repeats this result using the full sample from 2014 to 2017. We find a smaller overall effect when the initial years of data are included. This is supported by Figure 2, which illustrates the result visually, showing the average monthly variable profits of multi-store and single-store firms throughout the sample period. They begin the sample roughly equal and multi-store firms slowly gain a profit advantage which becomes significant in 2016 and continues to grow through the end of the sample period. This shows that the profit advantage is not simply a feature of an initial period when the industry was just ramping up. Instead, the advantage associated with being a member of a chain is growing with time. This figure also illustrates the magnitude of the effect of firm size on store profits.

Columns 3 and 4 compare the effect size for firms in their first year of operation relative to the period after the first year. The profit advantage associated with being part of a multi-store firm is

substantially larger later on in a store’s life compared with a smaller advantage in the first year. Again this suggests the advantage is not temporary but is something that grows over time. Finally, column 5 shows the result when the dependent variable is log of profits, to account for potentially large variance and skewness in this variable.

Appendix 8 shows this causal effect of scale on profits is robust to a series of potentially confounding factors. These include robustness tests for strategic entry into markets with less competition for lottery slots, replication of the results with product fixed effects included, and tests for the impact of store advertising.



2.1.1 Entry Timing Tests:

Next we explore the profit advantage enjoyed by multi-store firms by studying how the timing of entry by the stores in a multi-store chain effect their partner firms’ profits. One possible explanation for the profit advantage caused by being in a multi-store firm is that firms can use the profits generated by their first store to open a larger or higher quality second or third store, thus generating

Table 6: Timing of Main Effects

	(1) All Multi-Store Firms	(2) First Entrants Only
Entry order=2 or 3	-58124.2 (54960.972)	
# Stores open>1		19495.6 (16068.771)
# Applications FE	Yes	Yes
Store Age in Months FE	Yes	Yes
Firm Age in Months FE	Yes	
Month-Year FE	Yes	Yes
Market FE	Yes	Yes
Observations	324	207
R^2	0.806	0.808

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: This table shows results for two subsamples of multi-store firms. The dependent variable in both columns is variable profits. Column 1 shows results for all stores in multi-store firms and shows the effects of being the 2nd or 3rd entrant on variable profits. Column 2 shows results only for stores who were the first entrant in a multi-store firm and shows the effects of the 2nd firm's entry on the focal stores variable profits.

higher average variable profits. Table 6 shows the result investigating this mechanism. Column 1 of Table 6 shows that, among multi-store firms only, stores that enter second or third are not more profitable compared to the first store opened by that firm. In fact, after conditioning on firm age and time fixed effects, later entrants are less profitable than first entrants but not statistically significantly so. The lower profits experience by these stores may be due to a strategic decision to open in more desirable locations first and less desirable locations later. This rules out a story where the multi-stores are more profitable per store through a larger investment in fixed costs in second or third stores after successfully operating one store for some period of time.

Similarly, we compare the effect of being in a multi-store firm on the first entrants in each firm by using the timing of future store entries. Column 2 of Table 6 uses only first entrant members of multi-store firms and shows that when the number of stores in the firm goes from 1 to 2, the profits of the first entrant increase conditional on age, time and market fixed effects. Again, the result is not statistically significant due to the small sample size. This is suggestive that the profits of the first entrant increase after the second entrant opens, though.

2.2 Prices, Wholesale Costs, Sales, and Margins

Next we explore the mechanism by which larger scale causes higher profits. The theoretical literature on retail competition makes ambiguous predictions on whether larger or more prominent retailers will charge higher or lower prices than smaller retailers, all else equal. This makes this setting especially interesting to investigate the effect of scale on pricing decisions as a primary mechanism. This also serves as a test of to what extent greater market power leads retailers to increase prices. Two stores owned by the same firm should internalize the pricing externality and raise prices as a result.

We begin by using transaction level data on sales and prices to construct average retail prices, average wholesale prices, variable profits per unit sold, and total units sold. We use data at the product level to allow for product fixed effects, whereas variables were calculated at the store level in the previous subsection. We define product at the manufacturer-category level, thus if a manufacturer makes multiple products in the same category we aggregate these together.

Table 7 shows the effect of being in a multi-store firm on each of these variables. The effect on variable profits is smaller because the unit of observation is now product as opposed to store. We observe that, as expected, multi-store firms have lower average wholesale prices than single-store firms. This is the traditional notion of economies of scale in retail, and results from volume discounts and greater bargaining power with suppliers.

We also observe that multi-store firms set significantly lower retail prices than single-store firms for the same products. This difference does not merely result from the multi-store firms passing along the lower wholesale costs to final retail prices. Wholesale prices are lower by \$.21 on average, compared to a \$.36 difference in retail prices. Consequently, multi-store firms have significantly lower variable profits per unit sold than single-store firms. They are able to have significantly higher overall variable profits despite this because they sell substantially more total units. Multi-store firms sell 54 more units of each product per month than single-store firms, a roughly 14% higher sales volume.

It is notable that the profit advantage enjoyed by stores that are members of multi-store firms comes through lower prices and lower margins and not the opposite. Several plausible advantages of chain membership would be expected to result in higher and not lower prices. These include including greater reputation or consumer awareness, lower competition since chain stores do not compete with one another, and higher quality stores with larger assortments. The effect of a

Table 7: Effect of Multi-Store Firm Membership on Primary Outcomes

	(1)	(2)	(3)	(4)	(5)
	Variable Profit	Retail Price	Wholesale Price	Profit/Unit	Units sold
Multi-Store	154.1*** (36.183)	-0.36*** (0.082)	-0.21*** (0.039)	-0.15** (0.054)	54.1*** (11.654)
# Applications	Yes	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes
Market FE	Yes	Yes	Yes	Yes	Yes
Product FE	Yes	Yes	Yes	Yes	Yes
Observations	173618	173618	173618	173618	173618
R^2	0.231	0.726	0.751	0.608	0.201

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: This table shows the effects of membership in a multi-store firm on 5 outcome variables. In each column the dependent variable is described in the column header. Data sample period is April 2016 to April 2017. The observation level is product-store-week.

reduction on competition by removing a competitor should be even more valuable in this setting due to the lack of potential entry. We observe that the median number of stores in a market is 4 and stores that are part of multi-store firms tend to be located near each other, with a median distance of 16 miles and 90% of stores are within 20 miles of their chain partners. Therefore these multi-store chains have potentially substantial pricing power in their local areas and yet still charge lower prices than stores operating alone.

We note that removing the product fixed effects from these specifications yields nearly identical results for retail and wholesale prices, and slightly larger effects on units sold and thus variable profit. This suggests that the results are not driven by differences in composition of products offered, where multi-store chains offer more low quality products.

While the previous result focused on the last year of the data when store entry and pricing had stabilized, in Table 8 we show how the results for each of these dependent variables changes over the full 2014 to 2017 sample period. Notably, we see a clear time trend in that the difference in average prices charged by multi-store firms is growing over time. The price difference is negligible in 2015 but grows to \$.69 by 2017. This represents a substantial gap in average prices as the median price of 1 unit is \$6.65 in 2017. We see no clear trend over time in wholesale prices, where there is a clear advantage to multi-store firms in 2015 but no change over time. This suggests the growth in the pricing gap is not caused by simply passing along to customers the wholesale costs savings

Table 8: Change Over Time in Effects of Multi-Store Firm Membership

	(1) Variable Profit	(2) Retail Price	(3) Wholesale Price	(4) Profit/Unit	(5) Units sold
Multi-Store=1 \times year=2015	-427.1*** (79.616)	0.071 (0.168)	-0.28*** (0.080)	0.35** (0.111)	-89.4*** (22.172)
Multi-Store=1 \times year=2016	122.5** (41.471)	-0.20* (0.087)	-0.22*** (0.042)	0.013 (0.058)	33.4** (11.549)
Multi-Store=1 \times year=2017	202.4** (61.627)	-0.69*** (0.130)	-0.25*** (0.062)	-0.44*** (0.086)	68.7*** (17.162)
# Applications	Yes	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes
Market FE	Yes	Yes	Yes	Yes	Yes
Product FE	Yes	Yes	Yes	Yes	Yes
Obs	251867	251867	251867	251867	251867
R^2	0.219	0.718	0.732	0.604	0.204

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: This table shows the effects of membership in a multi-store firm on 5 outcome variables. In each column the dependent variable is described in the column header. The observation level is product-store-week.

associated with chain membership.

Price margins, as measured by variable profit per unit sold, also fall substantially over time, from a positive \$.35 gap in 2015 to negative \$.44 in 2017. At the same time, there is a large relative increase in the number of units sold per product, from 90 fewer units per month in 2015 to 69 more units in 2017.


This growth in the price and profit advantage associated with multi-store chain membership over time are consistent with the overall pattern shown in Figure 2. We visualize the difference in retail prices over time in a similar way in Figure 3, and for wholesale prices in Figure 4.

2.3 Assortment

This section investigates results for assortment decisions. Choosing how many and which products to stock are key decisions of a retailer. The theoretical literature on retail strategy and customer preferences has grown to emphasize that the number of products sold and consumer search over multiple products have significant interactions.

Retailers in this industry carry products from a set of five main categories, but carry potentially

Figure 3: Retail Price Gap Grows Over Time




Note: This figure shows average retail prices over time for multi-store firms and single-store firms. The grey area represents 95% confidence intervals.

very large numbers of varieties within those categories.¹³ Carrying products from a manufacturer incurs a fixed cost, and multi-store chains may be able to split that fixed cost over multiple stores such that the benefit of carrying one more product outweighs the cost for a larger number of products. We investigate whether this is true following the same empirical strategy as in the previous section.

Using transaction level data we construct the number of unique products sold in each store in each month as well as the number of brands or manufacturers purchased from in each month. Results are shown in Table 9, which shows the effect of being a member of a multi-store chain on assortment size defined these ways. Multi-store firms offer larger assortments defined in both ways compared to single-store firms, on a per-store basis. Figure 5 also demonstrates this visually. At the beginning of the sample single-store and multi-store firms offer similar assortment sizes but

¹³The categories are: usable marijuana (leaf) products, solid edible products, liquid products, extract products for vaporizers, and “other” miscellaneous products. Together these account for over 95% of product sales and products are defined by these categories in the regulatory transactions data.

Figure 4: Wholesale Price Over Time




Note: This figure shows average wholesale prices over time for multi-store firms and single-store firms. The grey area represents 95% confidence intervals.

over time the multi-store firms increase their per-store assortments. The disparity in number of products and number of brands offered increases consistently over the course of the sample.

These combined results showing multi-store firms offer larger assortments and charge lower prices are consistent with the model of Rhodes (2015), which shows that when consumers desire multiple products and have search costs, firms that offer more products will attract a larger but more price-sensitive set of consumers. This model makes specific predictions about consumer demand which we explore in the next section. An alternative explanation is that larger firms such as multi-store chains have greater name recognition or awareness in a market and thus may be the first choice for consumers searching for a specific product or low price. In this case, Armstrong et al. (2009) predicts that this prominence gives these stores an incentive to charge lower prices to deter customers from searching elsewhere. We test the prediction of this model in Appendix 7 and do not find support for prominence effects causing lower prices.

Figure 5: Number of Unique Products and Brands Over Time



3 Consumer Demand

In this section we estimate store-level consumer demand elasticities for each firm type over time. The previous section has showed that multi-store firms eventually charge significantly lower prices and margins than single-store firms, offer larger product varieties, and earn significantly higher profits. This leaves open the question of whether this strategy is replicable by single-store firms. That is, does each firm type face the same consumer demand and or do multi-store firms face a different marginal customer in terms of price sensitivity? Answering this will also further shed light on what mechanism results in the substantially higher profits earned by multi-store chains.

We proceed to separately estimate the store-level elasticity of demand faced by each firm type over time. We follow Hoch et al. (1995) and Hitsch et al. (2019) in estimating and aggregating demand elasticities using category-level prices and sales volumes. Let \mathcal{J}_{st} consist of the categories sold at store s at week t . We construct these at the weekly level and estimate the following log-log demand equation:

$$\log(q_{jst}) = \alpha_s + \sum_{k \in \mathcal{J}_{st}} \beta_{jk} \log(p_{jst}) + \tau_t + \epsilon_{jst} \quad (1)$$

We construct sales-weighted average prices p_{jst} at the category level. α_s are store fixed effects and τ_t are county-week fixed effects. These account for local time-varying demand shocks and store-level time-invariant factors such as location or market-level demographics. We estimate this equation

Table 9: Multi-Store Firms Offer Larger Assortments

	(1) Number Unique Products	(2) Number Brands
Multi-Store	6.82** (2.451)	5.15*** (0.932)
# Applications	Yes	Yes
Age FE	Yes	Yes
Time FE	Yes	Yes
Market FE	Yes	Yes
Obs	4637	4650
R^2	0.126	0.465

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: This table shows the effects of membership in a multi-store firm on 2 measures of assortment size. In each column the dependent variable is described in the column header. Number of products is the monthly number of unique inventory IDs at the store level. Number of brands is the monthly number of processors purchased from at the store level.

separately for each category. The result are estimates of average own-price elasticities β_{jj} and cross-price elasticities β_{jk} for every category and category-pair.

After including category, county-week, and store fixed effects there remains the possibility that there are unobserved demand shocks at the store-week level that are correlated with prices and cause price to be endogenous. This would bias price elasticity estimates towards zero. We use the availability of product-level wholesale prices to construct cost-shifting instruments. In particular, we construct the average category-level wholesale prices of all stores in a market and use these as instruments for retail prices. This assumes that wholesale prices have a direct effect on retail prices but not on final demand. The reason we use market-level averages is that if wholesalers observe the same store-week demand shocks and have sufficient market power to adjust store-specific wholesale prices, those prices would also respond to the same demand shocks. Using market-level wholesale prices avoids this potential issue.

3.1 Results of Demand Estimation

The estimated own-price elasticities are negative and lower than 1 for all categories. We thus focus on the usable (leaf) marijuana product category which accounts for over 84% of sales. We estimate demand separately for different firm type and time periods. Of particular interest is whether multi-store firms face different marginal demand than single-store firms. Because we observe that

multi-store firms charge lower prices and earn higher sales, if they face the same demand curve as single-store firms our estimated price coefficients might be closer to 0 as they are simply pricing in a more inelastic region of the same demand curve. In other words, we cannot conclude the two firm types face different demand if the multi-store firms have price elasticities closer to 0 in absolute value since this might just indicate movement along the demand curve rather than different demand curves.

Table 10 shows the estimated store-level elasticities by store type. Elasticities are estimated with the prices of the other categories included in the regression but their coefficients are excluded for space. We see that the OLS results are close to 1 for both firm types. When instruments are included elasticities are substantially more negative, indicating the OLS results are biased towards zero. The multi-store price coefficient is significantly more negative than the single-store coefficient.

Table 10: Estimated Price Elasticities

	(1)	(2)	(3)
	OLS	IV	IV
Single-Store \times ln(Price)	-1.01*** (0.032)	-1.75*** (0.157)	
Multi-Store \times ln(Price)	-1.08*** (0.110)	-3.79*** (0.284)	
Single-Store \times Year=2015 \times ln(Price)			-0.64*** (0.041)
Single-Store \times Year=2016 \times ln(Price)			-1.22*** (0.041)
Single-Store \times Year=2017 \times ln(Price)			-1.67*** (0.059)
Multi-Store \times Year=2015 \times ln(Price)			0.0041 (0.129)
Multi-Store \times Year=2016 \times ln(Price)			-1.24*** (0.145)
Multi-Store \times Year=2017 \times ln(Price)			-2.92*** (0.361)
Store FE	Yes	Yes	Yes
Market*Time FE	Yes	Yes	Yes
Observations	23068	23068	23068
R^2	0.882	0.866	0.891

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: This table shows estimated price elasticities for each store type in different subsamples based on year. In each column the dependent variable is ln(Units sold). Both units sold and price are for the usable marijuana product category. Elasticities are estimated with the prices of the other categories included in the regression but their coefficients are excluded for space.

Next, we test how these elasticities vary over time. Given that the stores and products are new and there is rapid entry in 2014 and 2015, we expect consumer demand responses may shift over time. We modify the demand model in equation 1 to allow for separate β_{jk} coefficients for each store type each year in the data.

We find very inelastic demand in the beginning of the data, perhaps as few stores have entered

and firms are still learning the about demand. Elasticities are below 1 in absolute value indicating firms may charge higher than the profit-maximizing price. In 2016 the coefficients for both types are more elastic, and multi-store firms' demand is similar to single-store firms'. By 2017 demand is more elastic for both store types, presumably due to greater competition stores who were awarded licenses in the lottery continue to enter the industry during the full sample. By 2017, multi-store firms face substantially more elastic demand than single-store firms. The pairwise comparison of 2017 coefficients is highly significant.

This difference in elasticities is true despite the fact that multi-store firms are also charging significantly lower retail prices than single-store firms. These results suggest that over time consumer demand shifts outwards for multi-store firms, as evidenced by their higher profits and sales, but also that the marginal consumer visiting a multi-store firm is substantially more price sensitive than the marginal consumer faced by a single-store firm.

4 Market Concentration and Markups


In this section we investigate the relationship between market concentration and markups. There is a growing interest in whether or not there is a long-term trend towards firms charging higher markups as shown by (De Loecker and Eeckhout (2018) and De Loecker, Eeckhout, and Unger (2018)). Some speculate that this increase in markups is caused by greater industry concentration as a result of lax antitrust policy or the growth of superstar firms. We explore this mechanism in our setting, which has exogenous variation in firm size and market concentration and where we can observe markups directly. Previously, we have shown that larger firms charge lower prices and lower markups. The analysis so far has been conducted at the store level. We now turn to the market as the unit of analysis and measure the effects of greater market concentration on average market-level prices and markups.

Conditional on the number of licenses allocated to a market and the number of applications filed in the lottery, whether or not a market contains stores that are part of a multi-store chain is effectively the result of random chance. Here, concentration should be measured at the firm level not the store level. This is because stores that are owned by the same firm will internalize the effects of price competition among themselves, and the presence of multi-store chains in a market correspond directly to higher market concentration.

We first show the negative relationship between concentration and markups holds in the cross section of markets in our data. To do so we calculate the level of concentration in a market as the normalized Herfindahl-Hirschman Index, a measure of concentration that adjusts for the number of firms in a market. It is calculated as:

$$H^* = \frac{H - \frac{1}{N}}{1 - \frac{1}{N}} \quad (2)$$

where $H = \sum_i^N s_i^2$, s_i is the market share of firm i , and N is the number of stores in a market. In Figure 6 we show the significant and positive relationship between this measure of concentration and the proportion of stores in a market that are members of a multi-store chain, excluding markets with none.



Note: This figure shows the relationship between the share of stores in market that are part of multi-store chains and the normalized Herfindahl-Hirschman Index index. Only markets with at least one chain store are shown.

Then, we measure the causal effect of this increase in concentration on markups and prices. To do so we estimate a simple matching model comparing markets with multi-store chains to those without. Our approach consists of matching markets using the number of licenses allocated and the number of applications filed, both of which are pre-lottery outcomes. We argue that conditional on

these two variables, assignment of the treatment is effectively random. We perform the match using nearest neighbor matching, calculated using the Mahalanobis distance. This algorithm is convenient for this case because we match on two integer-valued variables and less transparent procedures like propensity-score matching are not needed.

As in the previous section, we use data from 2016-2017 to calculate average prices and markups weighted by sales at the market-month level. Table 11 shows results for the matching estimation as well as an OLS specification. The OLS specification includes fixed effects for the number of licenses, the number of applications filed, and month.

Table 11: Effects of Concentration on Market Outcomes

	(1)	(2)	(3)	(4)
	OLS	OLS	Matching	Matching
	Avg. Price	Avg. Markups	Avg. Price	Avg. Markups
Presence of Multi-Store Firm(s)	-0.59*** (0.035)	-0.058*** (0.008)	-0.69*** (0.028)	-0.041*** (0.006)
Number of Licenses FE	Yes	Yes		
Number of Applications FE	Yes	Yes		
Time FE	Yes	Yes		
Observations	3367	3367	3367	3367
R^2	0.728	0.638		

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: This table shows the results from OLS and matching regression of the effect of greater concentration on average market level prices and markups. Concentration is measured using a dummy for the presence of multi-store firms and the unit of observation is a city-month. In Columns 3-4 matching is done using a nearest neighbor algorithm using number of licenses and number of applications to match, with exact matching on month.

Our results show that both prices and markups are significantly lower in markets with multi-store chains, despite the fact that these markets have higher concentration. These effects are large in magnitude and robust to specification.

Combined with the previous results showing that larger firms charge lower prices, these results

provide evidence against some stylized models of firm concentration and markups. Under these models, firms compete in Cournot outcomes and greater concentration mechanically translates to higher markups. Our results show that when many strategic variables are available to firms and are chosen simultaneously, and when consumers respond endogenously, large firms can increase their profits substantially in ways other than charging higher markups than small firms. Similarly, more concentrated markets may have lower average prices and markups than less concentrated markets depending on the endogenous choices of larger firms. Together, these results highlight the importance of using industry specific data with careful examination of the mechanisms involved to advance the debate on concentration and markups.

5 Discussion and Conclusion

We study the nature of economies of scale in the retail sector using a natural experiment that randomly allocated different numbers of retail outlets to different firms. We find that scale, in the form of number of outlets, has a significant impact on retailer strategy and firm outcomes. Larger retailers charge substantially lower margins, face lower wholesale costs, and earn higher total sales. The result is a profit advantage worth more than \$380,000 per store per year.

A simple comparison of retailer profits and chains size in this data would observe a strong positive correlation between number of outlets and profits, and therefore might conclude that better firms have higher profits and also grow and open more outlets. While this simultaneity effect is likely a true feature of retail chain growth, our data show there is also a direct effect of greater scale on profits. We can also conclude that the source of these economies of scale is not just the traditional cost advantage associated with large retailers. While this advantage is present, retailers do not simply pass their lower costs along to consumers but also charge significantly lower margins as well. The result is substantially lower prices at multi-store firms compared to stores operating alone.

It is not obvious *ex ante* whether larger firms should charge higher or lower prices, or whether they would face more or less price-sensitive consumers. Larger firms face less competition and more concentrated markets, which in isolation should allow them to charge higher prices. In addition, by having multiple outlets they may have greater consumer reputation or by offering larger assortments they offer a higher quality shopping experience. This might attract the less price-sensitive higher valuation consumers, leaving smaller single-store firms to compete for low valuation consumers by

offering lower prices.

Another view is that stores with greater awareness and larger varieties will attract a more mainstream consumer base, who will have lower average willingness-to-pay than stores that sell fewer products. This mechanism is demonstrated in Rhodes (2015), who shows how in the presence of search costs and consumers who value multiple products, large firms will carry a larger assortment and attract a more price-sensitive “mass market” set of consumers while small firms offer a more niche assortment and sell to high value consumers. Our evidence is consistent with this model. We find a large outward shift in demand associated with stores being members of a multi-store chain, but also a shift towards more price-sensitive consumers. The first effect is large enough to generate a substantial profit advantage caused by operating in a chain, however.

At the market level, our evidence is also informative on the debate over the causes of the long term increase in markups observed by De Loecker et al. (2020) and others. It has been widely speculated that this increase is caused by growing industry concentration as a result of lax regulation. While we do not claim to rule out this explanation, our results show that simple models producing positive relationships between concentration and markups are probably missing some important aspects of how firms compete in many industries. This suggests that more industry-level studies are needed where researchers can carefully account for differences in the nature of competition and consumer demand and can study causal effects in a clear and credible way.

Finally, while our results are limited to a set of relatively small players unlike truly large retailers like Amazon or Walmart, there is value in focusing on smaller firms for studying economies of scale. First, we would expect larger effects as firms vary between 1 and 2 stores vs 100 and 101, and generally if scale effects have decreasing returns it would be hard to replicate this type of study for very large retail chains. Second, even in large and mature industries the debate on retail concentration and retail mergers is often centered around cases with just 2 or 3 firms in a market, such as the Staples and Office Depot merger. And third, outcomes for small firms and entrepreneurs are often seen as especially important. Our results shed particular light on the value of scale to small entrepreneurs and the role of luck in determining entrepreneurial outcomes. A cross-sectional analysis of the entrepreneurs who entered this industry would find large differences in profits and number of outlets and again might assume that higher quality entrepreneurs were able to run their businesses more successfully and expand as a result, when the true causation in this case runs in the opposite direction. Therefore barriers to scale, such as capital constraints, legal red-tape, etc

are likely to be a significant impediment to entrepreneurial success if they keep firms too small.

Ultimately consumers may also benefit from a more concentrated retail sector with fewer but larger firms. This is clear from the lower prices, larger assortments, and higher overall sales associated with chain stores. Despite reducing competition, margins decrease rather than increase at these stores. This suggests that barriers to scale, including Washington's cap on the number of licenses a firm can own, ultimately decrease the benefits of the retail sector for firms and consumers alike.

References

- ANDERSON, S. AND A. DE PALMA (2006): “Market Performance with Multiproduct Firms,” *Journal of Industrial Organization*, 14, 95–109.
- ARMSTRONG, M. AND J. VICKERS (2020): “Patterns of Price Competition and the Structure of Consumer Choice,” *working paper*.
- ARMSTRONG, M., J. VICKERS, AND J. ZHOU (2009): “Prominence and consumer search,” *RAND Journal of Economics*, 40, 209–233.
- AUTOR, D., D. DORN, L. F. KATZ, C. PATTERSON, AND J. VAN REENEN (2020): “The fall of the labor share and the rise of superstar firms,” *The Quarterly Journal of Economics*, 135, 645–709.
- BARRETTE, E., G. GOWRISANKARAN, AND R. TOWN (2020): “Countervailing Market Power and Hospital Competition,” NBER Working Papers 27005, National Bureau of Economic Research, Inc.
- BERRY, S., M. GAYNOR, AND F. SCOTT MORTON (2019): “Do increasing markups matter? lessons from empirical industrial organization,” *Journal of Economic Perspectives*, 33, 44–68.
- BRONNENBERG, B. J. AND P. B. ELLICKSON (2015): “Adolescence and the Path to Maturity in Global Retail,” *Journal of Economic Perspectives*, 29, 113–34.
- DE LOECKER, J. AND J. EECKHOUT (2018): “Global market power,” Tech. rep., National Bureau of Economic Research.
- DE LOECKER, J., J. EECKHOUT, AND G. UNGER (2020): “The rise of market power and the macroeconomic implications,” *The Quarterly Journal of Economics*, 135, 561–644.
- ELLICKSON, P., S. HOUGHTON, AND C. TIMMINS (2013): “Estimating Network Economies in Retail Chains: A Revealed Preference Approach,” *RAND Journal of Economics*, 169–193.
- FOSTER, L., J. HALTIWANGER, S. KLIMEK, C. KRIZAN, AND S. OHLMACHER (2016): “The Evolution of National Retail Chains: How We Got Here,” *Handbook on the Economics of Retailing and Distribution*, 7–37.
- FOSTER, L., J. HALTIWANGER, AND C. KRIZAN (2006): “Market Selection, Reallocation, and Restructuring in the U.S. Retail Trade Sector in the 1990s,” *Review of Economics and Statistics*, 748–758.
- GOLDMANIS, M., A. HORTAÇSU, C. SYVERSON, AND Ö. EMRE (2010): “E-commerce and the Market Structure of Retail Industries,” *The Economic Journal*, 120, 651–682.
- GRULLON, G., Y. LARKIN, AND R. MICHAELY (2019): “Are US Industries Becoming More Concentrated?,” *Review of Finance*, 23, 697–743.
- HALL, W., D. STJEPANOVIC, J. CAULKINS, M. LYNKEY, J. LEUNG, G. CAMPBELL, AND L. DEGENHARDT (2019): “Public health implications of legalising the production and sale of cannabis for medicinal and recreational use,” *The Lancet*, 394, 1580–1590.

- HITSCH, G., A. HORTACSU, AND X. LIN (2019): “Prices and Promotions in U.S. Retail Markets: Evidence from Big Data,” *University of Chicago, Becker Friedman Institute for Economics Working Paper No. 2019-117*.
- HOCH, S., A. MONTGOMERY, B. KIM, AND P. ROSSI (1995): “Determinants of Store-Level Price Elasticity,” *Journal of Marketing Research*, 32.
- HOLLENBECK, B. (2017): “The Economic Advantages of Chain Affiliation,” *RAND Journal of Economics*, 48.
- HOLLENBECK, B. AND K. UETAKE (2019): “Taxation and Market Power in the Legal Marijuana Industry,” *Working Paper*.
- HOLMES, T. (2011): “The Diffusion of Wal-Mart and Economics of Density,” *Econometrica*, 79, 253–302.
- HORTACSU, A. AND C. SYVERSON (2007): “Cementing Relationships: Vertical Integration, Foreclosure, Productivity, and Prices,” *Journal of Political Economy*, 115, 250–301.
- HOSKEN, D. AND S. TENN (2016): “Horizontal merger analysis in retail markets,” *Handbook on the Economics of Retailing and Distribution*, 250–286.
- ILANES, G. AND S. MOSHARY (2019): “Market Structure and Product Assortment: Evidence from a Natural Experiment in Liquor Licensure,” *Kilts Center for Marketing at Chicago Booth â Nielsen Dataset Paper Series 2-013*.
- INDERST, R. AND G. SHAFFER (2007): “Retail Mergers, Buyer Power and Product Variety,” *The Economic Journal*.
- KIMBALL, M. S. (1995): “The Quantitative Analytics of the Basic Neomonetarist Model,” *Journal of Money, Credit and Banking*, 27, 1241–1277.
- KLOPACK, B. (2018): “One size fits all? The value of standardized retail chains,” *working paper*.
- LOERTSCHER, S. AND L. M. MARX (2019): “Countervailing Power,” .
- MCAFEE, R. P. (1994): “Endogenous Availability, Cartels, and Merger in an Equilibrium Price Dispersion,” *Journal of Economic Theory*, 62, 24–47.
- MORAGA-GONZALEZ, J. AND V. PETRIKAITE (2013): “Search costs, demand-side economies, and the incentives to merge under Bertrand competition,” *RAND Journal of Economics*, 44, 391–424.
- MRÁZOVÁ, M. AND J. P. NEARY (2017): “Not so demanding: Demand structure and firm behavior,” *American Economic Review*, 107, 3835–74.
- RATCHFORD, B. (2016): “Retail productivity,” *Handbook on the Economics of Retailing and Distribution*, 54–72.
- RHODES, A. (2015): “Multiproduct Retailing,” *Review of Economic Studies*, 82, 260–390.
- RHODES, A. AND J. ZHOU (2019): “Consumer Search and Retail Market Structure,” *Management Science*.

- SYMEONIDIS, G. (2010): “Downstream merger and welfare in a bilateral oligopoly,” *International Journal of Industrial Organization*, 28, 230–243.
- SYVERSON, C. (2004a): “Market structure and productivity: A concrete example,” *Journal of Political Economy*, 112, 1181–1222.
- (2004b): “Product substitutability and productivity dispersion,” *Review of Economics and Statistics*, 86, 534–550.
- (2019): “Macroeconomics and Market Power: Context, Implications, and Open Questions,” *Journal of Economic Perspectives*, 33, 23–43.
- TYNDALL, J. (2019): “Getting High and Low Prices: Marijuana Dispensaries and Home Values,” *Real Estate Economics*.

6 Data Appendix

6.1 Application Data

The list of businesses that have applied to licenses is available at the Washington State Liquor and Cannabis Board website.¹⁴ This list of licenses is not cumulative as we noticed that some licenses are dropped from the file through time. To recover the history of all license applications we use the website wayback machine. It allows us to recover all the listings made available to the public since the market opened. We use this procedure to recover the list of processors, producers, and retailers that have ever applied to a license. In 22 instances, firms receive a new license number but maintain their operation at the same location. We treat these cases as continuously operating firms.

6.2 Transaction Data

We have two distinct data sets that are put together to form the final transactions data.

- Retail dispensing data: contains all transactions between retailers and consumers with timestamp, prices, quantity, product type, strain, and parentid. The parentid variable indicates a 16 digit barcode identifier of the batch or lot the sample was taken from. It displays the company making the sale but it does not have the exact license that was responsible for the sale.
- Inventory transfers data: contains all transactions between the upstream and downstream markets. Importantly, it displays the information at the license level. Other variables that are included in this data are: strain, type, quantity, sale price, and parentid.

The parentid variable indicates a 16 digit barcode identifier of the batch or lot the sample was taken from. This variable is also present in dispensing and allows us to match the datasets above.

6.3 Outliers

As with any administrative data, the data contains a small fraction of errors, misentries, and outliers. We systematically delete observations believed to be mis-entered into the BioTrack system.¹⁵ Namely, cases where the final sales price is below \$3 per gram or above \$80 per gram (0.8% of transactions), wholesale prices below \$1 or above \$30 per gram (.04% of transactions), weight below .5 grams or above 30 grams (.07% of transactions) and markups above 3 (.04% of transactions).¹⁶

6.4 Taxation

We check for whether retailers enter tax-inclusive or pre-tax prices into the dataset. This first requires collecting sales tax rates for every store in every month because sales taxes may vary at the 9-digit zip code level. We find the 9-digit zip code of each store and match each store to the correct sales tax in each month of the data.

¹⁴<https://data.lcb.wa.gov>

¹⁵We follow the same procedure as Hollenbeck and Uetake (2019) for the removal of outliers and treatment of taxes, described below.

¹⁶Legal purchase limits are one ounce for usable, 16 ounces for solid, 72 ounces for liquid, and 7 grams for concentrates.

Since the majority of final prices use integer units, we check for the share of integers generated by each possible data entry rule. These rules include entering the pre-tax price, the price with excise and sales taxes included, and the prices that include either excise or sales taxes alone. Then at the retailer-month level we choose the rule that generates the highest share of integer prices, in some cases we also compare pricing within a retailer-category from month to month and checking final prices against the market average in each month to insure consistent treatment. We find that prior to the tax change in July 2015, roughly 8% of retailers enter tax-exclusive prices, 60% enter prices that include excise but not sales taxes, and 25% enter fully tax-inclusive prices. After the tax law changes, over 90% of retailers enter tax-exclusive prices. Once we recover the rule at the retailer-month level we construct the correct tax-inclusive and tax-exclusive prices for every transaction.

6.5 Advertising

We acquired data on retailer advertising from Kantar Media, a firm that tracks advertising spending across media. They have tracked advertising spending in the cannabis industry since before legalization through the current time for the Seattle and Spokane markets, including outer suburbs. We observe total spending at the monthly level for each media: outdoor, newspaper, magazines, television, radio, and internet display. Advertising in this industry is heavily regulated and restricted, and consequently there is relatively little of it.¹⁷

Table 12: Summary Statistics for Kantar Media Advertising Data

	Ad Spending (\$)	Share
Media		
Internet Display	2,094,816	11.5%
Local Magazines	52,922	.3%
Local Radio	404,485	2.2%
Newspapers	3,578,534	19.7%
Outdoor	12,028,456	66.2%
Total	18,159,213	
Market		
Seattle	16,996,795	93.5%
Spokane	1,162,418	6.4%

We use Kantar Media’s product description, which includes the store name, to match by hand the advertising data to the main data. Table 12 provides summary information on firm advertising. Almost all of it takes place in Seattle, broadly defined. Two thirds of all spending is on outdoor billboards, followed by roughly 20% spent on newspaper ads.

¹⁷For instance, all advertising is prohibited “in any manner that would be especially appealing to children or other persons under 21 years of age.” Advertising within 1000 feet of a school is also prohibited and all advertising must include a variety of cautionary text. Outdoor billboard advertising is limited to providing the name and location of a business.

6.6 Co-branded and Former Medical Stores

The regulation WAC 314-55-155, from May 2016, regulated the advertising in the WA cannabis industry. Among other rules, it allowed firms to have their own website where an “our story” tab is usually found. This feature and the wayback machine digital archive allowed us to map the lottery winners that had experience in the cannabis industry prior to the legalization as collective gardens. We explore whether stores formerly operating as medical dispensaries prior to I-502 have different likelihoods of being present in multi-store firms or whether including this as a covariate changes any of our results. We find that former medical stores have lower sales and profits consistently, even at the very beginning of the sample when they might have an advantage based on awareness and reputation. No other results are effected by including this distinction.

The same data features made it possible to identify sets of stores that operate under the same brand, which might encompass one or more firms. This is because their store locations were listed on their websites. We find several instances of co-branded stores that are part of separate firms, meaning they have separate owners and acquired their licenses through separately filed lottery applications but at some point they chose to align themselves in their marketing. They either share the same name or have a similar name and list all stores on a joint website.

An empirical concern in treating these as multi-store firms is that these arrangements may be more likely to be entered into by higher quality managers or they may be more likely for stores with prime locations. We therefore exclude co-branded but legally distinct sets of stores from our definition of multi-store.

7 Prominence and First Mover Effects

We run a test designed to help determine the mechanism behind our results from section 2. In the introduction we laid out two theories for why larger firms would charge lower margins than smaller firms. The first was that firms that offer more products attract a more price-sensitive set of consumers (Rhodes (2015)). The second was that larger firms such as multi-store chains have greater name recognition or awareness in a market and thus may be the first choice for consumers searching for a specific product or low price. In this case, Armstrong et al. (2009) predicts that this prominence gives these stores an incentive to charge lower prices to deter customers from searching elsewhere.

In section 2, we have shown that multi-store firms do offer larger assortments and in section 3 we show they face different consumer demand. Here, we want to measure whether the effects of being in a multi-store chain on prices is the result of higher prominence. To do so, we compare the multi-store effect to the effects of prominence as defined by the first store to enter in a market. Given the newness and novelty of this industry, the first store entering a market was likely to generate a high degree of awareness among interested customers. This may generate a first mover effect that acts similar to the proposed mechanism for multi-store firms.

Our results for prominence are displayed in Table 13. We test for a first mover prominence effect causing lower prices in the overall sample in column 2. Compared to the result for multi-store firms in column 1, we find no similar effect for first movers. We also test for effects in small markets in columns 3 and 4. Small markets are defined as those for which fewer than 10 retailers ever entered, and we speculate that a first mover effect would be stronger in these markets. Again, we find a stronger result for multi-store firms in these markets but no effect for first entrants. In columns 5 through 8 we repeat these tests for only the latter part of our sample, 2016 and 2017 and find similar results. A similar set of tests using variable profits as the dependent variable tell a consistent story, with multi-store firms earning significantly higher profits but first movers seeing no profit advantage in general, and significantly lower profits in small markets.

Altogether the, the lack of a first mover effect on prices provides evidence that multi-store firms are not led to set lower prices as a result of greater prominence in the style of Armstrong et al. (2009).

8 Robustness Tests

8.1 Robustness on Lottery

In this section we present a set of robustness checks for the primary results on variable profits. First, we show in Table 14 tests regarding the random distribution of licenses via the lottery. Conditional on filing a valid application, the result of the lottery is a uniform draw over all firms. Nevertheless, sophisticated firms might have strategically chosen markets based on their beliefs about the number of participants that would enter in each lottery. In this case, these firms would be more likely to win multiple licenses, albeit in less desirable markets.

In Table 14 we show how our main results on profits across stores change if we include a set of additional variables related to the lottery outcomes. Column 1 shows the baseline result. Column 2 shows results with a dummy indicator for lottery vs non-lottery licenses. Column 3 includes city fixed effects and shows the effect of including the probability of winning as a covariate, where probability of winning is calculated as the number of applications filed divided by the number of

Table 13: Test of Prominence Effects

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Markets:	All Markets		Small Markets		All Markets		Small Markets	
Time Periods	All	All	All	All	2016-17	2016-17	2016-17	2016-17
Multi-Store	-0.30*** (0.076)		-0.82*** (0.223)		-0.38*** (0.078)		-1.15*** (0.238)	
First Entrant		-0.022 (0.050)		0.063 (0.086)		-0.063 (0.056)		0.062 (0.092)
# Applications	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Market FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Product FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	251867	251867	97530	97530	207487	207487	79649	79649
R^2	0.718	0.718	0.728	0.728	0.725	0.725	0.732	0.732

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: This table shows the effects of multi-store chain membership and first mover status on prices. In each column the dependent variable is average retail price. First entrant a dummy that equals 1 for stores that were the first store to enter a given city. The observation level is product-store-month. Small markets are defined as those in which fewer than 10 retailers ever enter.

licenses available. Column 4 includes the number of applications at the market level directly. In column (3) we observe that effectively expanding the market definition has little effects on store profits. None of the other columns show a significant relationship between store profits and the additional variables, and the effect of multi-store firm membership on store profits is effectively unchanged.

Next we analyze whether the distribution of licenses allocated in the lottery is significantly different from a random allocation. While the licenses were distributed according to random draws, if firms were able to choose markets strategically there might be a non-random relationship between number of applications and number of licenses won. We therefore test whether the joint distribution of licenses is significantly different than what might occur if all licenses had the same win probability. To do so, we calculate the expected number of licenses won if they were distributed from a poisson process where all applications had an equal probability of winning a license. We use a poisson parameter λ equal to the overall probability determined by the total number of applications and licenses.

Table 15 shows the result of row-by-row chi-square goodness-of-fit tests for the joint distribution of applications and licenses under the assumption of independent poisson draws. For 3 out of 4 rows we cannot reject the null of independent draws with p-values between .3 and .4. For the row with 3 applications per firm, we do reject independent draws with $p = .038$. Essentially, among firms that applied for 3 licenses too few won multiple licenses. Combining all rows, for the sample as a whole we cannot reject independent draws, however.

Table 14: Robustness Checks on Lottery

	(1)	(2)	(3)	(4)
Multi-Store	31799.4*** (5512.789)	31482.0*** (5576.544)	25535.8*** (7400.843)	33184.3*** (7369.916)
All Lottery		1179.4 (3106.230)		
Pr(win)			117465.0*** (21485.698)	
# Applications (Lottery)				29.9 (47.953)
# Applications	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
County FE	Yes	Yes		
City FE			Yes	Yes
Obs	3262	3262	2187	2187
R^2	0.579	0.579	0.598	0.593

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: The dependent variable in each column is variable profits. Data sample period is April 2016 to April 2017.

Table 15: Test of Independence: Distribution of Applications and Licenses

	Applied	Approved			χ^2	p-value
		1	2	3		
1	117	0	0			
2	60	10	0	1.80	.406	
3	76	18	3	6.52	.038	
4	26	8	0	.85	.357	
5+	23	10	3	.83	.362	

Note: This table shows the joint distribution of applications filed and stores ultimately won in the 2014 retail lottery. We calculate the expected number of licenses in each bin assuming all applications have an equal probability of winning and calculate a chi-square goodness-of-fit test independently for each row.

8.2 Store Advertising

In this section we test whether the multi-store profit advantage is attributable to higher levels of advertising spending. First, we compute monthly ad spending by each store type and find that multi-store firms do not advertise more basis than single-store firms. Multi-store firms spend \$540 per month on ads on average, with a median spending of \$0 and a standard deviation of \$3238. Single-store firms spend \$734 per month on ads on average, with a median spending of \$0 and a

standard deviation of \$2849. The figures for multi-store firms are at the firm level since we cannot separately measure advertising at the store level.

Table 16 shows the primary results for profits when advertising is included as a covariate. For multi-store firms we show firm-level advertising and store-level advertising, calculated by dividing the firm level amount by the number of stores. We also include $\log(\text{adspend}_{st} + 1)$ due to the highly skewed nature of the data. In Columns 2-4 of Table 16 we see that advertising is positively correlated with variable profits, but including this covariate does not weaken the main effects.

Table 16: Robustness of Primary Outcomes to Advertising

	(1)	(2)	(3)	(4)
Multi-Store	31799.4*** (5512.789)	32704.7*** (5411.957)	34067.7*** (5404.094)	37616.0*** (5354.251)
Firm Ad Spending		3.52*** (0.322)		
Store Ad Spending			3.87*** (0.336)	
log(Ad Spending)				5315.7*** (368.713)
# Applications	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
Market FE	Yes	Yes	Yes	Yes
Observations	3262	3262	3262	3262
R^2	0.579	0.595	0.596	0.605

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: The dependent variable in each column is variable profits. Data sample period is April 2016 to April 2017.