Munich Personal RePEc Archive

Taxation and Market Power in the Legal Marijuana Industry

Hollenbeck, Brett and Uetake, Kosuke

UCLA Anderson School of Management, Yale School of Management

10 March 2021

Online at <https://mpra.ub.uni-muenchen.de/106618/>
MPRA Paper No. 106618, posted 16 Mar 2021 09:53 UTC

Taxation and Market Power in the Legal Marijuana Industry*

Brett Hollenbeck

Anderson School of Management

UCLA [†]

Kosuke Uetake

School of Management

Yale University[‡]

April 25, 2020

Abstract

In 2012 the state of Washington created a legal framework for production and retail sales of marijuana. Ten other U.S. states and Canada have followed. These states hope to generate tax revenue for their state budgets while limiting harms associated with marijuana sales and consumption. We use a unique administrative dataset containing all transactions in the history of the industry in Washington to evaluate the effectiveness of different tax and regulatory policies under consideration by policymakers and study the role of imperfect competition in determining these results. We use both a reduced form sufficient statistic approach and structural methods to show a number of results. First, Washington's strict cap on firm entry has resulted in retailers with substantial market power. This market power has immediate consequences for both state tax revenue and consumer welfare. Second, because these entry restrictions have caused retailers to behave like local monopolists, the state could substantially increase revenue generated from marijuana legalization by acting as the retailer itself, as it did for alcohol sales until 2012, without a large increase in prices. Third, despite having the nation's highest tax rate at 37%, marijuana in Washington is not overtaxed as many policymakers in other states have argued. The high taxes do not result in lower revenue or a substantial black market. Instead Washington is still on the upward sloping portion of the Laffer curve and the amount of revenue generated by a tax increase is significantly larger due to retailer market power than it would be under perfect competition. Our results suggest there is not widely available black market marijuana competing with legal retail sales. Finally, the high excise tax is primarily borne by consumers and not by firms, and there is a large social cost associated with each dollar raised.

Keywords: Legal marijuana, retail pricing, market power, tax incidence, pass-through, regulation

*The authors thank the Morrison Center for Marketing Research for generous funding. We wish to thank Renato Girollo for excellent research assistance. We also are grateful for comments from Sylvia Hristecheva, Fernando Luco, Eugenio Miravete, Keaton Miller, Matt Osborne, Peter Rossi, Stephan Seiler, and Katja Seim as well as seminar participants at UBC Sauder, UC Berkeley - Haas, Boston College, Lehigh University, University of Tokyo, the International Industrial Organization Conference 2017, Marketing Science 2017, and the Econometric Society 2018 Summer Meetings.

[†]brett.hollenbeck@anderson.ucla.edu

[‡]kosuke.uetake@yale.edu

1 Introduction

By June of 2019, eleven U.S. states had passed laws legalizing the purchase and sale of cannabis products for recreational use and are in various stages of creating and implementing regulatory systems for legal sales, production, and distribution of this product as well as its taxation. Once they have done so roughly 30% of the U.S. population will live in states with a legal retail cannabis industry. Canada has also passed a nationwide legalization in 2018. In 2019, this industry accounted for \$12.2 billion in the U.S. and \$14.9 billion worldwide.¹

Similarly to alcohol, states have chosen to tightly regulate this industry due to concerns over public health issues related to marijuana sales and consumption, particularly user health, impaired driving, use of the product by minors, and possible ties to criminal activity.² Much like when the prohibition of alcohol was ended, states that are developing rules for this new industry face a number of regulatory and policy decisions. They share the same stated policy goal, namely taking the production and sales of this product out of the shadows so that it can be monitored, shaped via regulation, and taxed to raise revenue. This revenue can then be used to provide public services or reduce taxes elsewhere.

Despite having similar objectives, the novelty of the industry and the competitive setting has resulted in significantly different choices made by different policymakers regarding basic questions including how and how much to tax sales at the retail and upstream levels and how to design the industry's market structure. State excise taxes on marijuana products range from 10% in Maine and Massachusetts to 37% in Washington, nearly 4 times higher. The stakes of these choices are large, as the consequence of a difference of this size for a large state amounts to hundreds of millions of dollars per year in revenue. We focus therefore on three policy questions. First, is Washington, with the nation's highest tax rate, maximizing revenue? Is marijuana instead overtaxed, leading to loss of state revenue and widespread black market consumption? This Laffer Curve effect is widely cited by U.S. states and Canadian provinces as the primary reason to keep tax rates low.³ Policymakers in Washington and many other states consider this an urgent debate, with legislation

¹Data on current sales and forecast for future marijuana sales growth come from Arcview Market Research and BDS Analytics.

²See, for example, Gavrilova, Kamada, and Zoutman (forthcoming) on the effect of legalizing medical marijuana.

³The Laffer curve, defined as the relationship between the tax rate and total revenue raised, is usually considered in a macroeconomic issue describing the relationship between income taxes and labor supply. A similar relationship should apply to any commodity taxes as well, as the tax pushes the price upwards ultimately reducing demand. We note also that, while Arthur Laffer popularized this relationship, as pointed out by Auerbach (1985), the concept should be originally credited to Dupuit (1844).

introduced in 2016 that would lower the tax rate and suggesting this would increase revenue and reduce black market sales.⁴

Second, what is the incidence of taxes in this industry? When retail sales and production were made legal, three groups stood to benefit: consumers, the new firms entering the industry, and the state government via enhanced revenues that can pay for additional public services or reduce taxes elsewhere. The extent to which the tax burden is borne by consumers versus producers, and the social costs of each dollar of revenue generated are of direct interest and shed light on this question.

Third, many U.S. states and Canadian provinces have strictly capped the number of entrants allowed in this industry. This decision helps the state monitor and control marijuana sales, but necessarily leads to reduced competition and greater firm market power. We therefore study the importance of firm market power and imperfect competition and highlight the role this plays in our results on tax incidence, on state tax revenue, and on total marijuana consumption. Standard models of tax policy in public finance generally rely on assumptions of perfectly competitive markets which are unlikely to hold in these types of settings.⁵ In particular, firms with high margins can and will strategically adjust these margins in response to any change in regulation or taxation. Capturing this strategic firm behavior is essential to evaluating any potential policy changes, as we show. States effectively have two goals, maximizing tax revenue and minimizing overall marijuana consumption and black market sales, and two policy levers, the regulation of firm entry and the tax rate. We contribute to the broader economic understanding of how these tools interact, and how when governments can influence the degree of market power via regulation this will impact tax policy objectives.

We are aided by an exceptionally rich and comprehensive new source of data. Washington state's tight regulatory regime led to the creation of administrative data containing all transactions ever conducted in the state, including prices. Notably, in addition to all retail transactions, we also observe all upstream transactions. This data goes back to the first legal sales in 2014 through the present. Observing upstream data at the transaction level in a setting with unregulated prices is unusual, and we take advantage of this feature to improve our analysis in a number of ways.

⁴HB 2347 was introduced in January 2016 and proposed lowering the 37% tax rate to 25% arguing that "Lowering the retail marijuana excise tax will result in more state tax revenue due to the increase in sales which will follow."

⁵As also noted in Miravete, Thurk, and Seim (2018a), in their textbook *Public Economics*, Atkinson and Stiglitz (2015) comment: "We went on to emphasize that the model underlying much of the Lectures - and much of public economics - was the Arrow-Debreu model of competitive general equilibrium. Looking back a third of a century later, we are struck that little seems to have changed in this respect."

First, we are able to directly observe retail margins at the product level. Retailer market power is central to our results and observing these margins lets us measure that market power in a direct way, rather than taking the traditional approach of imposing a structural model of firm behavior to estimate margins. Second, we can measure the pass-through of cost shocks to final retail prices in a transparent reduced-form way. As we discuss below, this pass-through rate can be used as a sufficient statistic for supply and demand elasticities that lets us calculate tax incidence directly. Third, when we estimate a model of consumer demand, we are able to use upstream transactions to calculate novel instruments to better identify price elasticities. Fourth, when we evaluate counterfactual regulatory and tax policies, we can use observed wholesale costs as inputs rather than estimates.

We use this data to answer our research questions using a combination of structural methods and reduced form sufficient statistics. We use a reduced form estimate of cost pass-through to directly infer tax incidence and the social cost of taxation. We then use a model of consumer demand to estimate price elasticities. These can be combined with observed margins to infer competitive conduct and we show how these are directly informative regarding whether the industry is on the upward or downward sloping region of the Laffer curve. Finally, to simulate a series of counterfactual regulatory and tax policies, we impose a model of supply-side competition and verify that it replicates observed pass-through rates. We focus the estimation and counterfactual analysis on the 2016-2017 period, after which the industry had largely stabilized in terms of prices and retailer entry and revenue.

Retail entry is heavily restricted, with a strict cap of 550 licenses to be awarded for retailers and retailers set very high margins, with an average retail price of \$13.6 per gram and an average wholesale price of \$4.7 per gram. These facts both imply that retailers have significant local market power. Monopolistic behavior is not a immutable feature of the marijuana industry but is instead a result of a policy decision to restrict entry. Monopoly power by retailers has important implications for tax policy, because firms with market power can strategically respond to any policy change by adjusting prices. Anderson, de Palma, and Kreider (2001a) show that the degree of monopoly power has a significant effect on the extent to which taxes will be passed through to consumers. For example, an increase in the tax rate will cause retailers to lower their margins to stay in a more elastic region of the demand curve, thereby bearing more of the tax change and causing revenue to increase at a faster rate than it would under perfect competition.

We use detailed retail transactions data to estimate a model of consumer demand for mari-

juana products in order to measure price elasticities. We employ demand estimation techniques for horizontally differentiated products developed in industrial organization to allow for flexible substitution patterns across products and for the marijuana category as a whole as prices or taxes change. Measuring the price elasticity correctly is crucial for understanding the effects of excise taxes on both revenue and consumption. We find demand elasticities for marijuana products are on average between 2.8–3.5. These are among the first structural estimates of demand elasticity for legal marijuana, and they suggest that demand for cannabis is similar to alcohol products, which has an elasticity in the range of 3–4.5, as opposed to tobacco products, which have an elasticity around .6–.7.⁶ The average elasticity for marijuana products in aggregate compared to the outside good is 1.04, significantly more inelastic than the elasticity for spirits.⁷ This result suggests there is not widely available black market marijuana for the marginal consumer. We show in section 5 that this elasticity also implies the industry is still on the upward sloping region of the Laffer curve.

Next, we use the data on production and wholesale prices to estimate the degree to which cost shocks are passed through to retail prices. A broad literature from trade to industrial organization has shown that cost pass-through is directly informative regarding firm market power and consumer demand.⁸ We find a pass-through rate significantly above 1 is robust to a variety of specifications. Pass-through greater than 1 is consistent with an industry with both high market power by retailers and highly log-convex demand, since as costs increase the retailer will face an increasingly inelastic marginal consumer.

We find three pieces of evidence suggesting there is not widely available black market marijuana available to consumers. First, the aggregate elasticity for marijuana products is close to 1. In other words, there is not significant substitution away from (or towards) legal marijuana as a whole when prices increase (or decrease). Second, our demand estimates imply that the vast majority of consumer substitution when prices change takes place across products within the same retailer. Consumers do not seem to shop around across various retailers with respect to prices, indicating they value their relationship with individual retailers. Third, pass-through significantly above 1 is not possible if consumers have a lower-priced substitute available that would prevent retailers from increasing prices even beyond the increase in their costs. These features along with very

⁶See Gordon and Sun (2015) or Becker, Grossman, and Murphy (1991) for estimates of cigarette price elasticity and Miller and Weinberg (2017) or Miravete, Thurk, and Seim (2018a) for estimates of beer and liquor products, respectively.

⁷Miravete, Thurk, and Seim (2018a) find an average category level elasticity of 2.8 for spirits.

⁸See, for example, Nakamura and Zerom (2010), Hong and Li (2017), Fabra and Reguant (2014), McShane, Chen, Anderson, and Simester (2016).

high observed retail margins all suggest retailers are not competing with black market sellers. The decision by Washington to cap the number of retailers and to close down unregulated medical dispensaries starting in 2015 likely explains this, and the failure to take these steps in states like California likely explains the continuing widespread presence of black market marijuana there.

Next, we use our results on pass-through to measure the incidence of taxes in this new industry as well as the social cost of taxation.⁹ We take the framework suggested by Fabinger and Weyl (2013) who show how firm pass-through can be used as a sufficient statistic to characterize the degree of market power and curvature of demand when calculating tax incidence. The advantage of the sufficient statistic approach is that the estimation is transparent and credible but leads directly to welfare conclusions. We find that taxes are borne primarily by consumers, with 34% falling on producers and the remaining 66% by consumers. These taxes effectively raise revenue but they also produce an unusually large social cost. We find that for a given dollar of increased tax revenue 2.4 dollars of combined producer and consumer surplus are lost. This large social costs arises principally because retailers have such a high degree of market power and because marijuana demand is fairly inelastic and highly log-convex.

Given estimates of demand and pass-through and a model of retailer competition, we can analyze a series of counterfactual tax and regulatory policies and show how state revenue, total marijuana consumption, and consumer surplus differ under them. We first show that a simple model of Nash-Bertrand price competition between retailers replicates our reduced form results on pass-through. Next, using the data in 2017 when the legalized market is stabilized, we show that if the state monopolized retail sales, as some states do for alcohol sales and some jurisdictions are considering for marijuana, prices would change only slightly. This is because the cap on retailer entry already produces monopolistic conduct by retailers. But the state could capture the revenue associated with retail sales. Retailer variable profits are \$594 million per year, almost twice as large as annual tax revenue.¹⁰ Alternatively, the state could allow more entry to increase retail competition. We find that greater competition between retailers would significantly lower prices, increasing both total marijuana consumption and tax revenue.

We next evaluate counterfactual tax rates and find that despite having the nation's highest tax

⁹Note that the social cost here does not incorporate externalities such as the effect of marijuana consumption on crime. Existing evidence on the externalities is still limited. For example, Hansen, Miller, and Weber (2020) finds no effect of the legalization of recreational marijuana on traffic fatalities in Washington and Colorado.

¹⁰At current tax rates, marijuana taxes already raised 1.4% of Washington's state budget in 2017. With the additional revenue a state system of retailers would raise this could have increased to 3.9% with no change in tax rates.

rate for marijuana products at 37%, Washington is still on the upward sloping region of the Laffer curve and significantly more revenue could still be raised with a higher tax rate. On the other hand, if Washington set taxes at 15% like many other large states, annual revenue would be lower by \$165 million, or roughly 50%. A simple extrapolation of this result to California, a state that taxes at 15%, implies that California will miss out on over \$800 million in annual revenue by undertaxing marijuana relative to Washington’s current 37% rate. Retailer market power plays a significant role in this result. We compare the change in revenue when retailers strategically adjust prices following a tax change to those where retailers act as price-takers and find that the change in revenue is 16% larger due to retailer market power. Overall, our counterfactual simulations imply that the taxes have a high social cost but marijuana is not overtaxed from a revenue perspective; raising the tax rate could still substantially increase tax revenue and modestly reduce marijuana consumption.

2 Relationship to Literature

This paper is related to several literatures. The first is the recent empirical literature on regulation and taxation of sin products. A recent line of work focuses on alcohol taxation and regulation using differentiated product demand estimates and models of oligopoly competition (see for instance, Waldfogel and Seim (2013), Miravete, Thurk, and Seim (2018a), Miravete, Thurk, and Seim (2018b), Conlon and Rao (2019), Aguirregabiria, Ershov, and Suzuki (2015)). The most notable of these is Miravete, Thurk, and Seim (2018a), who also examines a Laffer Curve under imperfect competition. Their setting is the Pennsylvania liquor market, where the state imposes a uniform markup rule upstream and monopolizes retail sales. They show that strategic behavior by alcohol distillers in setting prices significantly effects the shape and location of the Laffer Curve. We find a similar result in our setting where there is no government regulation of prices both in the retail market and in the wholesale market, and market power resides primarily with retailers. Whereas they find Pennsylvania is on the wrong side of the Laffer Curve, we find that Washington is still on the upward sloping region. Moreover, our detailed wholesale price data allow us to estimate retail pass-through, which is a useful summary statistic to derive welfare implications of the tax policy as proposed by Fabinger and Weyl (2013).

Other recent work studies excise taxes on sugar and sugar-sweetened beverages, focusing on the incidence of these taxes and to what extent they are passed-through to final retail prices. These

products have also been singled-out by policymakers for excise taxes due to their effects on consumer health. These include Khan, Misra, and Singh (2016), Cawley and Frisvold (2017), Seiler, Tuchman, and Yao (2018), Bollinger and Sexton (2018) among others. These studies generally find less than complete pass-through of taxes to retail prices. Allcott, Lockwood, and Taubinsky (2019) studies the interaction between different motives for taxing sugary drinks and suggests an optimal tax-rate on sugar-sweetened beverages.

There is also a new and growing literature on legal and illegal cannabis industries. Many papers focus on the effect of legalizing the marijuana market on criminal activities. Adda, McConnell, and Rasul (2014), for example, argue that the decriminalizing marijuana allows the police to focus other types of offenses not on drug-related crimes, and hence legalizing marijuana can reduce crime rate. Hao and Cowan (2017) study the spillover effects of recreational marijuana legalization in Colorado and Washington on neighboring states on marijuana-related arrests. They find the increase in marijuana possession arrests in border counties of neighboring states but no impact on juvenile marijuana possession arrests.

The papers close to ours are Jacobi and Sovinsky (2016), Hansen, Miller, and Weber (2018), Miller and Seo (2019), and Thomas (2018). Jacobi and Sovinsky (2016) use the survey data on (illegal) marijuana usage and accessibility to marijuana in Australia to estimate the demand for marijuana separately from its accessibility. They predict the Australian government could raise \$12 billion from the tax by legalizing the marijuana market. Hansen, Miller, and Weber (2018) study the effects of the change in tax structure in Washington in 2015 using the event study analysis and present results on the effects of tax reform on the vertical integration incentives and the short-term effects of the change on prices. They find that the tax scheme before the reform strongly encouraged the vertical integration and the reform increased the retail price by 2.3%. They estimate the retail pass-through of the increased excise tax, instead of wholesale prices like ours, and find that the excise tax pass through is quite heterogeneous. Miller and Seo (2019) study the extent to which alcohol, tobacco, and marijuana are substitute products using an Almost Ideal demand system and point out the implications of this substitution for total tax revenue following marijuana legalization. They find that marijuana legalization decreased alcohol demand by 5% but that Washington is on the upward sloping portion of the laffer curve. Thomas (2018) studies the welfare implications of license quota by estimating a cannabis demand with a supply side retailer entry model. She finds that allowing free entry raises the total surplus by 21% relative to the current quota system. Our

paper examines the role of market power on tax revenues and consumption by comparing alternative regulatory regimes such as the state monopoly and the uniform markup rule that have been used in the liquor market in the U.S.

Lastly, this paper contributes to the extensive empirical literature on pass-through. The literature is too lengthy to summarize fully here, but of particular relevance includes the papers on pass-through of sales taxes (see, e.g., Marion and Muehlegger, 2011; Conlon and Rao (2019)) and input prices (see, e.g., Dube and Gupta (2008); Nakamura and Zerom (2010)). In addition are empirical applications that use pass-through to study welfare issues in regulated markets, including those following the framework described in Fabinger and Weyl (2013). This includes Miller, Osborne, and Sheu (2017), who use data on the Portland cement industry and a similar framework to study the incidence of environmental regulations. Atkin and Donaldson (2015) use pass-through to study costs related to trade. Agarwal, Chomsisengphet, Mahoney, and Stroebel (2014) use the pass-through rate of airline fuel on consumer prices to study the welfare effects of fees in the airline industry.

3 Data and Industry Background

3.1 Regulation and Taxation

Our data come from the Washington State Liquor and Cannabis Board (WSLCB), the regulatory body that oversees the retail cannabis market. A November 2012 popular referendum was approved by Washington state voters 56 percent to 44 and led to the creation of this industry. The referendum directed the state legislature to create a set of regulations allowing the industry to develop and to generate revenue for the state. The state subsequently instituted I-502 creating a licensing scheme under the WSLCB. The state allows sales for adults age 21 or over and bars public use of the product, driving under the influence, or transporting the product outside the state. Counties and cities have the option of “opting out” of the system and maintaining a prohibition on marijuana in their jurisdictions. It remains illegal statewide to grow the plant at home without a license and the state continues to arrest and prosecute illegal growers.

By law, there are three types of firms licensed to enter the industry: retailers, processors and producers, distinguished by their position in the vertical structure of the industry and each with a

separate license. Processors and producers may hold both licenses, meaning vertical integration is allowed upstream but is barred for retailers.¹¹ Sellers must maintain health and safety standards, including the regular testing of their products in state-approved laboratories.

Federal guidelines issued by the Department of Justice require the state to take measures preventing the product from being sold outside the state, particularly into neighboring states where the product is not legal. Consequently, Washington requires all cannabis sales to be entered into a tracking system called BioTrack beginning when a seed is planted and following it to the final retail sale.

The database contains all transactions in the industry dating back to the first sales in November 2014. This includes the prices and quantities of all sales between producers and processors, processors and retailers, and retailers and consumers. This paper uses data spanning the period between November 2014 and September 2017 and amounting to roughly 80 million transactions worth \$2.5 billion. The data identify the firms involved in each transaction but contain no data that identify customers or give customer characteristics. Products are identified by their category, which will be described in more detail in the next section, as well as a brief written description in some cases.

The state initially capped the number of retail licenses it would grant at 334, with this number allocated at the county level.¹² The number was somewhat arbitrarily chosen to match the number of state liquor store licenses granted under the states historical Liquor Control Board, and were distributed across counties approximately according to population. The number of firms applying for retail licenses far exceeded the number of available licenses in most counties and the licenses were thus awarded via a lottery run in April 2014. In January 2016 the state expanded the number of licenses from 334 to 556 and simultaneously acted to shut down any remaining retailers operating illegally that had been holdovers from the pre-2014 medical marijuana industry, which had been largely unregulated.

Production licenses were available in three tiers corresponding to different amounts of square footage. The total square footage available for production was initially capped at 2 million then later raised to 8 million. Like in the retail space, far more firms applied for production licenses than were allowed under this cap, and so production licenses were also awarded via lottery. There is no

¹¹The state of Colorado passed a similar referendum in November of 2012, but that state set up regulations which require retailers to be vertically integrated with producers.

¹²There is an application fee of \$250 for a retailer license. Also, the licensee needs to pay a renewal fee every year. The renewal fee was \$1,062 in 2015, and \$1,381 after 2018.

Table 1: Marijuana Excise Tax by States

State	Sales Tax	Annual Revenue (2018)	Revenue per capita (2018)
AK	\$50/oz	\$20.6 million	\$27.9
CA	15%	\$209 million	\$5.3
CO	15% retail, 15% wholesale	\$244 million	\$42.8
MA	10.75%		
ME	10%		
NV	10% retail, 15% wholesale	\$88 million	\$29.0
OR	17%	\$110 million	\$26.2
WA	37%	\$369 million	\$49.0

¹ Other: CA \$9.25/oz flowers & \$2.75/oz leaves. ME levies various production taxes based on product type.

² Some localities also impose their own excise taxes.

³ Source: (Davis, Hill, and Phillips 2019), revenue estimated for some states.

limit on the number of processing licenses.

Initially the state levied a 25% sales tax on all sales between producers and processors, processors and retailers, and on the final sale. Thus, if the firms were not vertically integrated upstream each product would be taxed three times. This created a strong incentive for upstream firms to vertically integrate to avoid one layer of taxes, disadvantaging non-integrated firms.¹³ To remove this disadvantage and simplify the tax system, Washington changed the tax rate in July 2015 to a single 37% tax on final retail sales by value. The new rate was chosen to be revenue neutral when compared to the existing tax rates and to not affect the final retail prices.¹⁴

Table 1 reports the sales tax for 8 states that have already started the legalized cannabis industry.¹⁵ As the table shows, the sales tax rate varies significantly across states, ranging from roughly 10% in several states to 37% in WA. Washington charges the highest sales taxes on marijuana by a large margin.¹⁶

¹³See Hansen, Miller, and Weber (2018) for more description of the July 2015 tax change and its effects on vertical integration incentives.

¹⁴Because the tax change was designed to be neutral with respect to final retail prices as well as state revenue, we choose not to use this change to try to measure how retail prices respond to changes in tax rates. Attempting to do so would also be complicated by the fact that the tax change coincided with several other changes in the market, including closing down previously unregulated medical marijuana dispensaries. Finally, the tax change occurred relative early in the industry's history when prices were changing rapidly and firms were still entering. We focus most of our analysis on 2017 and the latter half of 2016 when the market had reached a more stable and mature state.

¹⁵Recreational use of marijuana became legal in Vermont in July 2018, in Michigan in December 2018, and will be legal in Illinois in January 2020.

¹⁶Washington also charges the highest liquor taxes in the U.S., at 20.5% plus a unit tax of \$3.7708 per liter. This corresponds to a 61.8% tax on a 1.75 liter bottle with a listed price of \$15.99. Washington also charges the 3rd highest tax on cigarettes at \$3.025 per pack of 20 cigarettes.

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

3.2 Descriptive Results

In this section we describe the key features of the data that motivate our empirical analysis. Appendix A explains in greater detail the data sources and how the data were cleaned and matched.

First, because it was initially advantageous for tax reasons to vertically integrate, and because the act of “processing” is relatively simple for the basic product, most producers applied for and received processing licenses. Consequently the majority of upstream firms are vertically integrated. Because there is very little actual processing for this product, the primary result of this integration is that the industry avoids upstream double marginalization. For processors who make edible products or other more exotic products, the share which are vertically integrated is much lower since the processing of those products is significantly more complex. In 2017, 93% of wholesale goods are sold by vertically integrated processors.

The term cannabis is used generally to refer to any products containing the active ingredient contained in the cannabis plant. This comes in several distinct forms. These are “usable marijuana”, which is the flower of the plant and is meant to be smoked directly, solid edible products, liquid edible products, and extract of the active ingredient meant for inhalation as vapor. These account for 96% of sales, with the remaining 4% consisting of a large number of niche products which will

Figure 2: Average Wholesale Price By Category Over Time (\$/gram)

largely be excluded from analysis.

Figure 1 plots the average (tax inclusive) retail price over 4 years in our data. Generally, retail prices decrease over time for all categories, particularly in 2014 and 2015. Since 2016, prices plateau. The figure shows that in 2017, the average retail price across all products was \$13.62 per unit including taxes, where 1 unit is either 1 gram or a standard product unit. We plot average wholesale prices over time by product category in Figure 2. Similarly to retail prices, wholesale prices decrease over time for all categories, but the average wholesale price paid by retailers (\$4.74 per unit) was much lower than the average retail price.

Based on the retail and wholesale prices, we find that retailers earn substantial margins, which we plot in Figure 3. The average markup on 1 gram of usable marijuana is \$3.15 out of a total tax-exclusive retail price of \$6.28, yielding an average margin of .50 for usable products and .52 for all products. Aggregating at the level of product type, retailer margins ranged from .33 to .67 with most retailers setting margins between .5 and .6. These margins are substantially higher than typical margins in retail settings. The median margin for U.S. grocery products, for example, has been estimated at roughly .3 (Hottman (2018)) with higher estimates of .45 in the U.K. (Thomassen, Smith, Seiler, and Schiraldi (2017)) and with an upper bound of .52.

Table 2: Price Summary Statistics (2017)

	Total Sales (grams)	Retail Price		Wholesale Price	
		Mean	Std Dev	Mean	Std Dev
All Combined Products	5,052.9	13.62	9.68	4.74	3.45
Usable Marijuana	3,585.7	9.02	2.89	3.13	1.27
Solid Edible	400.2	19.44	10.98	6.59	3.62
Liquid Edible	101.2	24.07	12.44	8.36	4.22
Extract	726.3	27.79	10.51	10.00	3.31
Other Products	239.0	12.86	9.09	4.39	3.13

Note: This table presents total monthly sales and average prices for each product type during the year 2017. Retail prices are tax-inclusive. Total sales refers to the average monthly total sales of all products in grams or the equivalent unit.

Figure 3: Histogram of Average Retail Margin

Figure 4: Average Monthly Sales by Type

Table 3: Firm Revenue in 2017

	# of Firms	Mean	5th Pctile	95th Pctile	Std	Max
Monthly Revenue (2017)						
Retailers	385	\$202,354.8	\$21,947.5	\$573,506.4	185,117	\$1,394,183
Processors	642	\$63,377.6	\$2,100.0	\$247,165.2	149,459	\$2,181,563
Producers	388	\$14,921.3	\$560	\$69,038.4	22,975	\$174,856

Note: This table presents summary statistics on the number of licensed firms of each type in 2017 as well as data on monthly revenues. Monthly revenue data are averaged over January-June 2017 at the firm level.

The strict cap on retail licenses and high margins suggest that retailers display a high degree of market power in their local markets and we observe that they capture most of the industry's revenues. Retailer revenue accounts for 66% of all combined revenue in the industry. Figure 4 and Table 3 show the average monthly sales of each firm type from the industry's creation. These illustrate how retailers capture an out-sized share of total industry revenue, averaging slightly more than \$200,000 in monthly revenue in 2017. There is wide dispersion in the level of sales at the retailer level however, with the 10 largest retailers averaging roughly \$1,000,000 in monthly sales.

By contrast, the upstream market is not particularly concentrated. Over 600 processors reported positive sales in September 2017, the final month of our data. The 10 largest processors accounted for 22.4% of those sales and the 50 largest processors accounted for just over half of all sales. While there are no restrictions on processor size, the upstream industry has yet to show signs of increasing

concentration.

Table 4: Transaction Summary Statistics

	Mean	std. dev.	min	max	50%
# of wholesalers per retailer	66.9	40.8	1	192	65
# of retailers per wholesaler	15.4	26.0	1	137	7
Wholesale market share per retailer	0.082	0.22	0.005	1	0.015
CV of wholesale price	0.087	0.141	0	0.742	0.017

Note: The table shows summary statistics of the transactions between retailers and processors. CV, the coefficient of variation, is calculated as the ratio of the standard deviation of prices over the mean price. We calculate the mean and the standard deviation for each product category, strain, and week.

Table 4 reports the summary statistics on the relationship between retailers and processors. The first row shows the number of processors with which each retailer has some transaction. On average, a retailer has 66.9 processors that it has purchased from at least once. By contrast, a wholesaler has about 15 retailers to transact, which is much smaller than the number of transacting wholesalers for a retailer. These facts would indicate that wholesalers may not have much bargaining power against retailers. The third row shows the share of sales from each processor per retailer. The average share is 8% and the median share is 1.5%. Since retailers have a lot of transaction partners, they are not much dependent on any particular processor. These facts indicate that different processors are close substitutes from the retailer’s perspective. Lastly, the bottom row shows the coefficient of variation, a common measure of relative price dispersion, for wholesale prices. The average CV is 0.08 and its median is 0.01. This indicates that processors do not change prices very much across retailers, or the processors’ market power is limited.

4 Empirical Framework

This section describes the empirical framework which will be used to study tax and regulatory policy. The 37% sales tax imposed by Washington is substantially higher than other sales taxes including excise taxes on products considered harmful such as alcohol and sugar.¹⁷ The state had several goals when setting such a high tax rate, primarily to generate revenue for the state and to keep prices high and thus consumption low and relatively contained. Other states with the same

¹⁷Washington imposes a 20.5% tax on the shelf price of alcohol in addition to a flat spirits liter tax of \$3.7708/liter. Beer faces an effective tax rate of 11%.

goals have nevertheless chosen very different tax rates and regulatory regimes.¹⁸ We seek to study the effectiveness of these taxes in raising revenue and suppressing consumption, as well as evaluating their incidence.

Because of the strict limits on entry imposed by the state and the high retail margins observed in the data, any analysis of these questions would be incomplete without accounting for the fact that firms have substantial market power. Many core results in regulatory and tax economics rely on assumptions of perfect competition. By contrast, Anderson, de Palma, and Kreider (2001a) show that under imperfect competition, taxes can be passed on to consumers more than fully. In an extension of this work, Fabinger and Weyl (2013) show how this result applies to a broad class of oligopoly settings and show how reduced form estimates of cost pass-through can be used in a straightforward way to estimate tax incidence, as well being a general tool to inform issues related to the effects of regulation on consumer and firm surplus. This framework has previously been applied empirically in Atkin and Donaldson (2015) and Miller, Osborne, and Sheu (2017), the latter of which we follow in certain respects.

The following section describes the theoretical framework for characterizing the effect of a change in tax rate on state tax revenue as well as the incidence of and deadweight loss from taxation. This framework requires detailed estimation of consumer demand and the rate of pass-through from costs to final retail prices. This section will describe the estimation of each of these in succession.

4.1 Pass-Through Rate

A key empirical measure of firm conduct is the degree to which cost or tax changes are transmitted to consumers in final retail prices. The pass-through rate is widely used both in public finance, as an indicator of tax incidence and deadweight loss, and in industrial organization as a measure of the degree of competition and differentiation in an industry. This is also the key input to the framework developed by Fabinger and Weyl (2013), which characterizes how tax incidence and the social cost of taxation under imperfect competition can be summarized by pass-through and a conduct parameter. In this section, we discuss estimation of the pass-through rate of wholesale prices for a range of different specifications.

Wholesale prices are typically estimated from the assumed supply-side first-order conditions as

¹⁸For instance, Maine and Massachusetts impose 10% tax rates. Alaska imposes no tax on retail sales but a \$50 per ounce tax on production, which amounts to just under 10% of the retail price.

marginal costs (see, e.g., Berry (1994)), but an advantage of our data is that we can directly observe them. Using these data, we first estimate the following model to obtain the pass-through rate.

$$p_{gjt} = \beta_0 + \beta_1 w_{gjt} + \beta_2 w_{-gt} + x'_{gjt} \beta_3 + \mu_g + \mu_j + \mu_t + \varepsilon_{gjt}, \quad (4.1)$$

where p_{gjt} is the per gram tax-inclusive monthly-average retail price by retailer g for product j in month t , w_{gjt} is the per-gram monthly-average wholesale price that retailer g pays for product j in month t , w_{-gt} is the average wholesale price that competitors pay in month t , x_{gjt} is a vector of variables for observed retailer-product characteristics such as THC content, μ_g is the retailer fixed effect, μ_j is the product-category fixed effect, and μ_t is the year-month fixed effect, which captures unobserved market-level heterogeneity and macro economic shocks. This specification is similar to the one used in the existing literature such as Miller, Osborne, and Sheu (2017), which estimate cost pass-through in the cement industry.

To estimate equation (4.1), we need to aggregate the original transaction-level data to the monthly-level. Aggregating at the monthly level is consistent with the previous section.¹⁹ We define “a product” by the combination of the category (i.e., {usable, solid edible, liquid edible, and extract}) and the strain, and calculate the average retail and wholesale prices for each month. Similarly, we calculate the monthly average THC content from state-mandated potency analysis. Standard errors are clustered at the city level.

Table 5 shows the results of the panel linear regression. The results show that own pass-through rates are significantly higher than 1 for all specifications. The coefficient on THC is positive and statistically significant. The competitors’ wholesale prices are positively associated with own prices, but the magnitude is smaller than the own pass-through and sometimes statistically insignificant. Thus, we find that cannabis retailers pass through their cost shocks more than perfectly.

Since transactions between processors and producers are subject to 25% tax before June 2015, the interpretation of the pass-through might be different. To deal with the concern, the specification in Model (4) uses the observations only after July 2015, and we find that the pass-through estimate is robust. In Model (5), we exclude observations from the counties that are located at the border between Oregon and Washington. Although the model has already controlled for retailer fixed effects, there might be different demand trends for those counties due to potential cross-state

¹⁹In the appendix, we present the results based on weekly observations and find the results are robust.

Table 5: Pass-through Estimates: Monthly Average Price

	(1)	(2)	(3)	(4)	(5)	(6)
Wholesale Price	1.715*** (0.0203)	1.640*** (0.0204)	1.648*** (0.0215)	1.641*** (0.0211)	1.649*** (0.0235)	1.654*** (0.0209)
THC		0.00763*** (0.00107)	0.00782*** (0.00116)	0.00989*** (0.00130)	0.00775*** (0.00120)	0.0104*** (0.00131)
Competitor Wholesale			0.260*** (0.0672)	0.158* (0.0619)	0.268*** (0.0728)	0.112 (0.0700)
Constant	5.592*** (0.147)	6.000*** (0.145)	3.984*** (0.548)	4.424*** (0.491)	3.910*** (0.597)	4.313*** (0.520)
N	330662	330652	265065	251689	240680	175788
R-squared	0.68	0.70	0.71	0.71	0.71	0.72

Note: The table reports pass-through estimates with monthly average prices based on the linear panel fixed-effect model. All models control for year, month, retailer, strain fixed effects. Standard errors are clustered at the city level. Model (4) uses the observations only after July 2015, when the tax policy changed. Model (5) uses the observations excluding counties that are at the border between Washington and Oregon. Model (6) uses the observations only after July 2016, when some medical licenses were converted to retail licenses. Significance levels: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

trafficking. Even after removing those states, we find that the pass-through is still robust. Finally, Model (6) uses the observations only after July 2016 when the number of retail licenses is increased. We find that the pass-through rate does not change much after the change in the market structure.

As discussed in Miravete, Thurk, and Seim (2018a) and Fabinger and Weyl (2013), pass-through greater than unity suggests the combination of high firm market power and highly curved or highly log-convex demand. The finding is consistent with other pass-through estimates that find evidence of pass-through rates greater than unity such as Miller, Osborne, and Sheu (2017) and Conlon and Rao (2019). In those studies, the authors find significant market power of retailers in the cement industry and the liquor industry, respectively.

To see the robustness of the previous results, we also estimate pass-through with another specification that deals with a concern that error terms might be auto-correlated. We take the first

Table 6: Pass-through Estimates: First Difference

	(1)	(2)	(3)	(4)	(5)	(6)
Δ Wholesale Price	1.527*** (0.0178)	1.524*** (0.0178)	1.528*** (0.0174)	1.565*** (0.0189)	1.537*** (0.0187)	1.573*** (0.0169)
Δ THC		0.00661*** (0.00126)	0.00662*** (0.00141)			
Δ Competitor Wholesale			-0.0801 (0.0493)	0.0595 (0.0434)	-0.0602 (0.0493)	-0.0293 (0.0293)
Constant	-0.193*** (0.00863)	-0.193*** (0.00861)	-0.194*** (0.0103)	-0.189*** (0.0113)	-0.189*** (0.00997)	-0.132*** (0.0102)
N	291748	291748	232567	223063	211084	158373
R-squared	0.26	0.26	0.26	0.26	0.26	0.27

Note: The table reports pass-through estimates with monthly average prices based on the linear panel fixed-effect model. All models control for year, month, retailer, strain fixed effects. Standard errors are clustered at the city level. Model (4) uses the observations only after July 2015, when the tax policy changed. Model (5) uses the observations excluding counties that are at the border between Washington and Oregon. Model (6) uses the observations only after July 2016, when some medical licenses were converted to retail licenses. Significance levels: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

difference of equation (4.1) to estimate the following specification.

$$\Delta p_{gjt} = \beta_0 + \beta_1 \Delta w_{gjt} + \beta_2 \Delta \bar{w}_{-gt} + \Delta x'_{gjt} \beta_3 + \Delta \mu_t + \Delta \epsilon_{gjt}, \quad (4.2)$$

where $\Delta p_{gjt} = p_{gjt} - p_{gjt-1}$, $\Delta w_{gjt} = w_{gjt} - w_{gjt-1}$. Other variables $\Delta \bar{w}_{-gt}$ and Δx_{gjt} are similarly defined. Note that retailer fixed effects and product fixed effects, μ_i and μ_j are all eliminated by taking a difference.

Table 6 reports the estimation results. Similar to Table 5, we find that the own pass-through estimates are still greater than unity for all specifications. The change in THC contents has a positive and statistically significant effect on the change in retail prices, but the magnitude is very small. The change in the average wholesale prices of competitors has insignificant effects.

To see how pass-through rates vary by product category, now we estimate the same specification

Table 7: Pass-through Estimates: First Difference by Type

	(1) Liquid	(2) Solid	(3) Extract	(4) Usable
Δ Wholesale Price	1.899*** (0.00590)	1.743*** (0.00931)	1.352*** (0.00588)	1.492*** (0.0153)
Constant	-0.185*** (0.00906)	-0.235*** (0.0141)	-0.345*** (0.0105)	-0.0813*** (0.0130)
N	47281	28319	92242	123906
R-squared	0.69	0.56	0.37	0.08

Note: The table reports pass-through estimates with monthly average prices based on the first-difference model by product category. All models control for year, month fixed effects. Standard errors are clustered at the city level. Significance levels: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

as equation 4.2 by category. Estimates reported in Table 7 show that the pass-through rate is greater than unity for all product categories. The pass-through is the highest for liquid products and the smallest for extract products.

Another concern one may have would be the fact that the recreational cannabis market in Washington is changing over time and the pass-through rates also vary month by month. We estimate the monthly pass-through and report the results in Figure 5. We find that the pass-through fluctuates a lot before June 2015 when the new tax policy is implemented. Since July 2015, the pass-through is stable, or slightly increasing, around 1.6.

As noted by Conlon and Rao (2019), one potential reason pass-through may be greater than unity is the use of discrete prices and discrete price changes. Conlon and Rao (2019) find that 77% of *quarterly* price changes in the distilled spirits market in Connecticut are in whole-dollar increments and it can lead to excessive pass-through for some types of tax changes. They suggest an ordered logit model of pass-through to capture this effect. We investigate the extent of discrete price changes in our data in Appendix B and find discrete pricing to be fairly common. We present results in the Appendix showing that an ordered logit model following Conlon and Rao (2019) produces similar pass-through rates near 1.5.

In the next subsection, we estimate a flexible model of consumer demand and simulate how firms would respond to small cost shocks under Nash-Bertrand pricing. We find a mean simulated

Figure 5: Pass-through rate by month

Notes: The figure plots the wholesale passthrough rate by month. The estimating model is based on equation (5).

pass-through of 1.43, suggesting that our measured pass-through is consistent with preferences and standard pricing behavior in this setting.

4.2 Demand Estimation

We next study consumer demand using a structural approach. We use a model of random coefficient nested logit (RCNL) demand in order to produce robust estimates of price elasticities. We use a model of demand that is nested at the retailer level to capture the retail structure of sales in this industry that form a key feature of both own and cross-price elasticities. RCNL demand is well-suited for settings like this when product differentiation takes a discrete form (e.g., Miller and Weinberg (2017)). Each city is a separate market and a “product” is defined at the retailer-category level. Following the discrete choice demand literature, we model demand over $j \in \mathcal{J}$ products in each market m in time period t for a set of consumers defined by i . Each consumer has utility which is modeled as

$$u_{ijt} = x_{jt}\beta_i^* + \alpha_i^* p_{jt} + \xi_{jt} + \bar{\epsilon}_{ijt}, \quad (4.3)$$

where x_{jt} is a vector of observed characteristics of both products and retailers and p_{jt} is the retail price. The observable product characteristics are product type and retailer intercepts. The term ξ_{jt} captures unobserved product quality that varies over product, market and time and is observed to firms and consumers but not the researcher.

To allow for heterogeneity in individual preferences, we model consumer utility over price as $[\alpha_i^*, \beta_i^*] = [\alpha, \beta]' + \Pi D_{im} + \Sigma v_{im}$, where D_{im} is demeaned consumer income and $v_{ij} \sim N(0, I_{n+1})$. The parameters α and β therefore capture the average effect of price and other characteristics on preferences, Σ captures the covariance in unobserved preferences over characteristics, and Π is a matrix of coefficients that capture the effect of and individual income on valuation of those characteristics. A specific form of heterogeneity we are interested in is varying preferences for marijuana by lower and higher income consumers. We therefore allow preference parameters to vary across the income distribution.

We follow Grigolon and Verboven (2014) in modeling correlation in preferences over certain products, in this case all products sold by the same retailer. This serves to capture the retail sector structure present in the industry and specifically it allows for the possibility of more substitution between products within a retailer than across retailers. The result is the random coefficient nested logit or RCNL model. Specifically, the idiosyncratic term $\bar{\epsilon}_{ijt}$ follows the nested logit distribution, where products in the same group have correlated preferences.²⁰ We can therefore write this term as:

$$\bar{\epsilon}_{ijt} = \zeta_{igt} + (1 - \rho)\epsilon_{ijt}, \quad (4.4)$$

where $\rho \in [0, 1]$ and represents a nesting parameter. The “nests” in this case are each retailer, as well as the outside good. As ρ goes to 1, products’ taste shocks become perfectly correlated within the store. Hence, any differences in demand across products within the same retail stores would thus be due entirely to observables entering the demand model. Plugging this expression into equation 4.3 gives

$$u_{ijt} = x_{jt}\beta_i^* + \alpha_i^*p_{jt} + \xi_{jt} + \sum_{g \in G} \chi(j \in g)\zeta_{igt} + (1 - \rho)\epsilon_{ijt}, \quad (4.5)$$

²⁰This has a similar effect as having random coefficients on preferences over retailers, but provides a more straightforward interpretation for the correlation of preferences over retailers and a closed form expression for the integral in the choice probability.

where $\chi(j \in g)$ is a dummy variable indicating if product j is in group g , meaning sold at retailer g . The full set of retailers in a market is noted by G . Allowing for a random coefficient on price and a flexible nesting parameter on product type allows for robust substitution patterns. When $\Sigma=0$, $\Pi=0$, and $\rho=0$, the model collapses to a standard logit demand.

The mean value of the outside option of not purchasing is normalized to zero. Defining the mean component of utility and the idiosyncratic components as:

$$\delta_{jt} = x_{jt}\beta + \alpha p_{jt} + \xi_{jt}, \quad (4.6)$$

$$\mu_{ijt} = (x_{jt}p_{jt})(\Pi D_i + \Sigma v_i). \quad (4.7)$$

This utility generates the following conditional probability that consumer i purchases product j from retailer g :

$$s_{igjt}(\delta_{jt}, \theta, v_i, D_i) = M \cdot \frac{\exp((\delta_{jt} + \mu_{ijt})/(1-\rho)) \exp(I_{igt})}{\exp(I_{igt}/(1-\rho)) \exp(I_{it})}, \quad (4.8)$$

where $\theta = (\beta, \alpha, \rho)$ and I_{igt} is an inclusive value term such that

$$I_{igt} = (1-\rho) \log \sum_{j \in G} \exp((\delta_{jt} + \mu_{ijt})/(1-\rho)) \text{ and} \quad (4.9)$$

$$I_{it} = \ln(1 + \sum_g \exp(I_{igt})). \quad (4.10)$$

Market is defined at the city level as the state determines the retail license cap at the city level, and within each market sales are aggregated at the monthly level. Next, we define product at the product type level for each retailer, where type is defined as either usable marijuana, solid edible, liquid edible, extract, or other. The model combines all sales of products within a category, thus averaging unobserved heterogeneity at the level of retailer-product each month.²¹ Retailer quality is addressed with retailer specific intercepts in the utility function. This allows for fixed factors like location and is interacted with time to allow for retailer quality to vary from month to month.

Prices are standardized to the price corresponding to 1 gram of each product. We then average sales and prices across all products of the same type sold at the same retailer in each month and use these to construct market shares. We ignore the potential for consumer stockpiling across months,

²¹In tests where the product is defined at the processor-retailer-type-month level to allow for potential brand effects, i.e. different preferences across processors, results come out largely the same.

in part because the product is largely perishable. We instead focus on long-term rather than short-term price responses by using monthly-level prices and market-shares. If there is significant consumer stockpiling it would cause us to overstate own-price elasticities.²²

4.3 Estimation and Identification

We estimate the model following the approach of Berry, Levinsohn, and Pakes (1995). We use a GMM estimator that interacts the structural demand side error $\omega(\theta)$ with a set of instruments Z , where the demand parameters are $\theta = (\alpha, \beta, \sigma, \rho, \Pi, \Sigma)$. Formally the GMM estimator is formed from the population moment condition $E[Z' \cdot \omega(\theta)] = 0$. The GMM estimate is

$$\hat{\theta} = \min_{\theta} \omega(\theta)' Z A^{-1} Z' \omega(\theta) \quad (4.11)$$

for some positive definite weighting matrix A . To construct the structural error $\omega(\theta)$ we use the modified BLP contraction mapping suggested by Grigolon and Verboven (2014) to obtain the unique vector $\delta^*(x_{jt}, S_{jt}, \theta)$, which maps the observed market shares S_{jt} into mean utility values. A 2SLS regression of $\delta^*(x_{jt}, S_{jt}, \theta)$ on product characteristics, price and various fixed effects with instruments Z then produces a residual term that is equivalent to $\omega(\theta)$. In our 2-step GMM we use $A = Z'Z$ in the first step and in the second step construct the heteroscedasity robust optimal weighting matrix clustered at the retailer level.

After including product type, time, retailer and market-time fixed effects in the model, there remains some unobserved component of utility ξ_{jt} which varies over time and within retailer and is known to firms when setting prices. The particular concern is a demand shock to a specific product type at a specific retailer at the monthly level. To deal with this endogeneity problem, we consider three types of instruments. Because we observe wholesale prices at the transaction level we are able to construct novel instruments to measure a variety of types of cost shocks that exogenously vary with final retail prices. These wholesale prices serve as a direct measure of marginal costs at the product level, but if upstream firms have market power, the wholesale prices may also be correlated with unobserved demand shocks appearing in utility. To avoid this but still take advantage of the upstream data, we construct instruments from the average of all wholesale prices of products of the same type from markets outside each of the focal market. The use of this instrument essentially

²²Another potential source of consumer dynamics would be addiction. Since we have no individual-level data, we do not specifically model consumer addiction to cannabis products.

assumes that co-movement in wholesale prices across markets are driven by cost shocks and not demand shocks after accounting for any statewide demand trends using time fixed effects.²³ To form these instruments, we construct 5 geographic regions in the state of Washington and calculate average wholesale prices at the type-month level for each region. Because these are constructed using wholesale prices, the relevant region is the region where each processor is located and therefore these instruments vary across retailers located in the same market who face different cost shocks based on which processors they purchase from.

We also observe prices further upstream from transactions between producers and processors. These prices reflect the wholesale market for whole plants, which are significantly more homogenous than the final products sold by processors to retailers. Producer prices are unlikely to be influenced by transitory demand shocks at the retailer-type level and therefore represent good cost-shifters for the industry as a whole. We construct average producer prices at the region-month level. These prices are linked to each retail transaction through the regional location of the processor of each product, so that two products of the same type sold by the same retailer might have different upstream prices if their processors are located in different regions.

Finally, because the raw product is an agricultural good and is grown outdoors in many cases, we use exogenous weather shocks as further cost-shifting instruments. Specifically, we collect data from the National Oceanic and Atmospheric Administration (NOAA) on average monthly rainfall and temperature at the county level and link this to the county locations of each producer. Again, we link these to final retail prices using the fact that we observe the full supply chain. We lag these variable one month and find they have a significant effect on retail prices after controlling for market-month fixed effects. Together, wholesale price instruments, producer prices, and weather shocks provide a substantial amount of exogenous variation in prices with which to identify price elasticities. In the next section we present and discuss first-stage results showing these instruments together are quite strong in terms of affecting final retail prices.

In addition to potential endogeneity of prices, Berry and Haile (2014) and others note that the heterogeneity terms introduce an additional identification problem into the estimation. In our RCNL specification, this means additional instruments are needed to ensure identification of Σ , Π ,

²³These are similar in nature to so-called “Hausman” instruments, which are widely used and are typically constructed using retail prices in other markets. Unlike retail prices, wholesale prices are likely more representative of costs and less likely to be correlated with the specific demand shocks making up the structural error. Also, in our setting, advertising by retailers/wholesalers are allowed, but very limited by regulations. Moreover, there is no national-level advertising.

and ρ , the nonlinear components of preferences. To identify these requires exogenous variation in the conditional shares of the inside goods, in this case the share of sales of product type j sold at a specific retailer. We use two types of instruments, the number of product types sold by the retailer in each month and the average values of the cost-shifting instruments described earlier for competing products within the retailer. The number of products is a standard instrument and is used by Miller and Weinberg (2017) among others. The average cost-shifters reflects variation in competing products marginal costs and should be correlated with the focal products market share and uncorrelated with the structural error.

We implement the demand estimation using the pyblp package and following best practices as described by Conlon and Gortmaker (2019), which we find to converge rapidly and consistently. This package makes it relatively straightforward to include approximations to the optimal IV in the sense of Chamberlain (1987) as described by Reynaert and Verboven (2014). To do so, we follow the procedure described in Conlon and Gortmaker (2019) in which we obtain an initial estimate of all parameters, solve for the structural errors $\hat{\xi}_{jt}$, and use these to construct Jacobian terms $\frac{\partial \hat{\xi}_{jt}}{\partial \rho}$, $\frac{\partial \hat{\xi}_{jt}}{\partial \Pi}$, and $\frac{\partial \hat{\xi}_{jt}}{\partial \Sigma}$. The initial fixed point iteration scheme uses the acceleration method of Varadhan and Roland (2008) and a Broyden-Fletcher-Goldfarb-Shanno optimization algorithm.

To allow for an outside good, we fix the size of a market as being 4 times the market population. This can be interpreting as allowing each resident of a market to purchase up to 4 grams of the product per month.²⁴ Standard errors are clustered at the retailer level. For income interactions we collect data from the 2010 Census. We use ZIP code level market definitions and take the distribution of income across 16 categories and take 100 draws from this distribution for each market observation.

4.4 Results of Demand Estimation

Results from this estimation are shown in Table 8. Results are presented for a simple logit demand and a range of specifications of RCNL demand with different interaction terms. In all cases price coefficients are negative and estimated precisely. In both cases the nesting parameter suggests a high correlation in preferences among products sold by the same retailer. This is consistent with

²⁴Different notions of market size have been tested and produce effectively identical results. As noted by Miller, Hansen, and Weber (2018) there is evidence of unusually high demand on the Washington-Oregon border in the period prior to when Oregon began legal recreational sales. These cross-border effects should be captured by our inclusion of market-month fixed effects. We also test estimation excluding these counties and present the results in Table 9.

high travel or search costs and results in much more substitution across products within a store than across stores in response to a price change. The interpretation of a very high nesting parameter is that most substitution in response to price changes takes place within retailers across products, rather than across retailers for the same product.

Across specifications, extract products and usable marijuana have the highest utility, with liquid edible products the least preferred category. Higher income consumers are less price sensitive than low income consumers, and have a lower overall preference for marijuana products and lower relative preference for the usable marijuana category.

Table 8 also shows how estimates of the price coefficient and average own-price elasticity vary across specifications. The median own-price elasticity in our preferred specification, shown in column (3), is -2.89. This is close to but slightly higher than the median own-price elasticity across consumer packaged goods of -2.62 (Hanssens (2015)). The simple logit demand model produces higher own-price elasticities and significantly higher aggregate elasticity, as would be expected due to the lack of retailer nesting. Figure 6 shows how own-price elasticity varies over time using estimated market-time fixed effects. Despite an increase in the number of retail stores over time, consumers grow more inelastic with respect to prices. This could be caused by the declining availability of black market marijuana as formerly unregulated medical dispensaries were closed or converted into legal retail stores.

We also calculate the total elasticity for the marijuana category as a whole relative to the outside good and show these in Table 8. For our preferred specification in Column (4), which includes all interactions and a random coefficient on the constant term, we find the category has an aggregate elasticity of -1.04 . This suggests most substitution takes place within the marijuana category with only very modest substitution to the outside good. By comparison, Miller and Weinberg (2017) find a category elasticity of -0.7 for retail beer. This stands in contrast to the liquor category, in which Miravete, Thurk, and Seim (2018a) find an aggregate elasticity of -2.8 . Policymakers in Washington and other states have expressed concern about the potential availability of black market products as a black market in sales to consumers would impede the states ability to both regulate the market and generate revenue. Because of the combination of high retail margins and high taxes, prices in the illegal market would almost certainly be significantly lower than in the legal market even in the absence of economies of scale in production costs. Nevertheless we find that demand is relatively inelastic for the category as a whole, suggesting there is not a widely available

Table 8: Demand Estimates

		(1) Logit	(2) RCNL-1	(3) RCNL-2
Price	α	-.154 (.003)	-.31 (.081)	-.063 (.026)
Usable Marijuana	β_1	3.10 (.038)	3.0 (.92)	.77 (.20)
Solid Edible	β_2	1.23 (.017)	.71 (.15)	.47 (.27)
Liquid Edible	β_3	1.01 (.038)	.49 (.11)	.36 (.19)
Extract	β_4	3.84 (.063)	2.03 (.43)	1.41 (.60)
Nesting Parameter	ρ		.63 (.12)	.62 (.20)
Income×Price	Π_1		.02 (.006)	.023 (.020)
Income×Usable	Π_2		-.12 (.074)	-.36 (.16)
Income×Constant	Π_3			-.81 (.074)
Random Coeff. on Constant	Σ_1			.66 (.58)
Median Own-Price Elasticity		-3.63	-2.85	-2.89
Aggregate Elasticity		-2.11	-1.13	-1.04
Type FE		Yes	Yes	Yes
Time FE		Yes	Yes	Yes
Retailer FE		Yes	Yes	Yes
Market*Time FE		Yes	Yes	Yes

Note: This table presents estimates of the demand system for different specifications. Product characteristics are price and dummies for type, date and retailer. IV estimation is done using GMM in each column with all 3 sets of IVs. In each column there are 31,502 observations at the type-retailer-month level coming from 2,727 markets where a market is city-month. Standard errors are robust and clustered at the retailer level.

Table 9: Demand Estimate Robustness

		(1) Seattle Only	(2) Exclude Seattle	(3) Exclude Oregon Border	(4) 2016-2017 Only
Price	α	-.15 (.036)	-.046 (.003)	-.16 (.052)	-.062 (.029)
Nesting Parameter	ρ	.60 (.06)	.58 (.12)	.61 (.03)	.61 (.11)
Own-Price Elasticity		-3.10	-2.51	-3.41	-2.80
Aggregate Elasticity		-1.53	-.71	-1.08	-.94
Type FE		Yes	Yes	Yes	Yes
Time FE		Yes	Yes	Yes	Yes
Retailer FE		Yes	Yes	Yes	Yes
Market*Time FE		Yes	Yes	Yes	Yes
N		14040	17466	28015	26945

Note: This table presents estimates of the demand system estimated on different sub-samples of the data. Each column uses the specification from column 4 of Table 8 with all income interactions and a random coefficient on the intercept. Note that due to differences in income between different subsamples and the presence of a price-income interaction term, the mean price coefficient would be expected to differ across subsamples. The number of observations is at the type-retailer-month level and markets are at the city-month level. Standard errors are robust and clustered at the retailer level.

black market where consumers may find substitute products.²⁵ That the marijuana category is fairly inelastic as a whole could also indicate the product is habit forming or addictive. If this is the case, there is nevertheless little evidence of a black market substitute available to supply the product outside the legal retail setting.

Finally, Table 9 shows results when the model is estimated on several relevant subsamples of the data. Columns 1 and 2 estimate the model using only greater Seattle and excluding greater Seattle. Seattle is the largest market and, including suburban outer Seattle, accounts for nearly half of retailer-month observations. The median own-price elasticity and aggregate elasticity are significantly higher in Seattle than in the more rural parts of Washington, likely due to the presence of more retail store options. Excluding Oregon border counties does not significantly affect the results, suggesting the interstate trafficking discussed in Miller, Hansen, and Weber (2018) is not an issue for estimating preferences. Finally, we estimate the model using only the years 2016 and

²⁵By black market we refer strictly to the presence of illegally produced marijuana available for purchase. There may still exist black markets for legally produced marijuana. In particular, legal purchases may be made by consumers over 21 and then later re-sold illegally to people under 21, presumably with an additional markup. In addition, legal purchases can be illegally taken out of state, a topic examined in Miller, Hansen, and Weber (2018). Moreover, a caveat of our argument is that we assume there are few consumers who buy on the black market exclusively. If the segment of consumers never substitute back and forth, we cannot say much about the black market from studying the legal market.

Figure 6: Own-Price Elasticity Over Time

Table 10: Price Endogeneity

		(1)	(2)	(3)	(4)	(5)	(6)	(7)
Price Coef.	α	-.054 (.001)	-.154 (.004)	-.159 (.015)	-.127 (.008)	-2.37 (14.77)	-.122 (.008)	-.152 (.004)
Cragg-Donald F-Stat			164.2	220.0	124.4	.07	207.6	186.9
Kleibergen-Paap F-Stat			24.9	32.1	24.3	.03	26.0	28.3
Wholesale Price IVs			X	X				X
Weather IVs			X		X		X	X
Producer Price IV			X			X	X	
Type FE		Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time FE		Yes	Yes	Yes	Yes	Yes	Yes	Yes
Retailer FE		Yes	Yes	Yes	Yes	Yes	Yes	Yes
Market*Time FE		Yes	Yes	Yes	Yes	Yes	Yes	Yes
N		31,504	31,504	31,504	31,504	31,504	31,504	31,504

Note: This table presents price coefficient estimates and first-stage test statistics for different combinations of price instrumentation strategies. Price coefficients come from a logit demand model. Standard errors are robust and clustered at the retailer level.

2017, after the mid-2015 tax change and after greater retail entry was permitted by the state. We find a lower median own-price elasticity and an aggregate elasticity just below 1. This is consistent with the pattern in Figure 6 and suggests that when the state closed formerly unregulated medical dispensaries this substantially eliminated the black market.

Table 10 explores how well our chosen instruments correct for potentially endogenous prices. It shows the estimated price parameter and first-stage test statistics for different combinations of instruments. Price parameters come from logit demand, and are more negative with instruments included, indicating price endogeneity is present in the data. Taken alone, the instruments composed of average wholesale prices outside the focal firm produce the highest F-statistic. The lagged weather instruments also have a significant effect on retail prices. The average producer-level prices appear to be quite weak and produce very low F-stats and insignificant estimates of mean price preferences.

Pass-Through In the previous section we found very high observed pass-through of wholesale prices. We now explore whether our estimated consumer demand model is consistent with those findings. Using the method described in detail in section 5.3, we can impose a standard supply-side model and simulate different outcomes assuming Nash-Bertrand pricing. We first solve for Nash-Bertrand prices at observed wholesale prices and then solve for a new set of Nash-Bertrand prices at wholesale prices plus a \$1 cost shock and take the difference. Figure 7 shows the distribution of outcomes at the retailer-category-month level. The mean simulated pass-through is 1.43. This suggests the observed mean pass-through of roughly 1.5 is consistent with standard pricing behavior and consumer preferences.

In sum, our pass-through estimates show that pass-through is greater than unity for all specifications, or retailers pass through costs to consumers more than 100%. This indicates that retailers enjoy a great deal of local market power and suggest that most of the tax burden falls on consumers rather than retailers. These results together strongly suggest the conclusion that there is not readily available black market marijuana acting as a substitute for legal marijuana sales. If this black market existed, retailers would not be able to pass-through their cost-shocks more than fully without losing excessive sales. This is subject to the same caveat that if there exist distinct consumer segments who purchase from the illegal and legal market and do not substitute between them, the legal market could still exhibit these features. In the next section, we employ the framework developed by Fabinger and Weyl (2013) to quantify the incidence of Washington’s excise taxes

Figure 7: Simulated Pass-Through

This figure shows simulated pass-through using the consumer demand estimated in section 4.3. Assuming firms engage in Nash-Bertrand prices we estimate equilibrium prices under observed wholesale prices and these prices plus a cost shock and take the difference. Each observation is a retailer-category-month.

by combining the pass-through estimates and the consumer demand estimates.

5 Policy Analysis

In this section, we use our empirical results to evaluate the current policy of the recreational cannabis industry in Washington. First, we calculate the incidence of Washington's 37% excise tax on marijuana as well as the social costs of these taxes. We do so using a sufficient statistic approach based around our estimates of cost pass-through. Second, we use a simple model of firm behavior and our estimates of price elasticity to see if Washington has set its tax rate at the optimal level and if the state is still on the upward sloping region of the Laffer curve. Finally, we impose a supply-side model of firm behavior and examine the regulatory policies on the supply side including a state monopoly on marijuana sales. We also use this model to evaluate how much additional tax revenue the state could earn with higher taxes, how much tax revenue the state would miss out on with a 15% tax, as well as what effect these would have on total consumption. In each case, we highlight the effect of retailer market power on these outcomes. Through these analyses, we aim to provide

clear guidance to policymakers on how to regulate the retail cannabis market.

5.1 Tax Incidence and Social Cost

The empirical results of the previous sections can be combined to evaluate the effectiveness of the state's regulatory regime along several additional dimensions, notably its effects on consumers and producers and the efficiency with which revenue is generated. We first adopt the framework of Fabinger and Weyl (2013) to show how firm pass-through can be used as a sufficient statistic for analyzing tax incidence and the social costs of taxation.

Spatial differentiation as well as the cap on retail licenses suggest potentially high levels of retailer market power, and accounting for this market power is important to properly measure the burden of taxation and how it is distributed between firms and consumers. Measuring this tax burden is of direct interest to policymakers and it can also inform us as to what extent each of three different groups are benefiting from the existence of the new marijuana industry: producers, consumers, or the state government via increased tax revenue.

To fix ideas, consider the effects of a unit tax under perfect competition. A tax of size t is applied such that $p_S = p_C - t$, where p_S is the price received by sellers and p_C is the price paid by consumers. In this case, the costs of this tax will be split between consumers and sellers, and the ratio of the marginal incidence of this tax paid by consumers ($\frac{dCS}{dt}$) to that paid by producers ($\frac{dPS}{dt}$) is $I = \frac{\lambda}{1-\lambda}$ where λ is the pass-through rate describing the effect of the tax on equilibrium price, i.e., $\frac{dp}{dt}$.

Under perfect competition, it is a classic result that this pass-through can be derived as:

$$\lambda = \frac{1}{1 + \frac{\epsilon_D}{\epsilon_S}}$$

where ϵ_D is the elasticity of demand and ϵ_S is the elasticity of supply. This provides the familiar result that the burden of a tax falls most heavily on the inelastic side of the market.

Fabinger and Weyl (2013) extend this principle to settings of monopoly and imperfect competition. Under a general model of symmetric imperfect competition, they show that the equilibrium can be characterized by

$$\frac{p - mc}{p} \epsilon_D = \theta, \tag{5.1}$$

where θ is a *conduct index* which summarizes the degree of competition in the industry and can be thought of as the ratio of actual margins to the margins that would be charged by a monopolist or set of firms colluding on the monopoly outcome. It thus ranges between 0 for perfect competition and 1 for monopoly. They go on to show that the marginal effect of taxation on producers is:

$$\frac{dPS}{dt} = -[1 - \lambda(1 - \theta)]q \quad (5.2)$$

and the marginal effect on consumers is

$$\frac{dCS}{dt} = -\lambda q$$

Thus, the tax incidence can be calculate as

$$I = \frac{\lambda}{1 - \lambda(1 - \theta)}, \quad (5.3)$$

where in this case

$$\lambda = \frac{1}{1 + \theta/\epsilon_\theta + (\epsilon_D - \theta)/\epsilon_S + \theta/\epsilon_{ms}}.$$

In oligopoly settings, pass-through now depends on ϵ_θ , the elasticity of conduct with respect to quantity, and ϵ_{ms} , the elasticity of marginal surplus, defined as $ms = p'q$. While these objects are difficult to estimate directly, under this framework we can instead substitute the reduced form estimate of pass-through to compute the tax incidence and dead-weight loss terms above. Pass-through therefore acts as a sufficient statistic for the nature of the competitive reaction to a tax change.

Calculating incidence still requires an estimate of θ , the conduct index. Rather than estimate θ as part of a larger structural estimation of demand function parameters and marginal costs, we take advantage of the fact that wholesale prices are observed and therefore retail margins are observed. We directly compared observed retail margins to the hypothetical margins that a monopolist would charge in order to estimate θ . We effectively calculate the hypothetical margins of a single monopolist using the elasticity of demand estimated in the previous section and equation 5.1. A more complete description of this counterfactual is described in the following section. We estimate an average $\hat{\theta} = .89$ with a median of .87, in other words observed margins are 89% of the hypothetical monopolist's margins. This is consistent with the demand estimates showing high preferences for

individual retailers and relatively little substitution across retailers, allowing retailers to behave as local monopolists.

Equation 5.3 gives the ratio of consumer harm to producer harm from a small unit tax increase. While in practice Washington uses ad valorem taxes on retail sales, in this section we evaluate the effects of a unit tax because this corresponds directly to our pass-through results. This allows us to measure the incidence of marijuana taxes in a straightforward way while imposing relatively few assumptions on the nature of competition. In the following section we use estimates of the demand function to evaluate potential changes in the ad valorem tax on retail sales. In addition, several other states including California do impose unit taxes.

Using our estimated $\lambda = 1.6$ and $\theta = .89$, implied incidence of taxes falls roughly 34% on producers and 66% on consumers.²⁶ We can directly derive from these equations the effect of a change in unit taxation using average total monthly sales of approximately 5,000kg in 2017. For a given \$1 increase in a unit tax, state revenue would increase by roughly \$5 million, consumer welfare would fall by the equivalent of \$8.0 million while producer profits would fall by \$4.1 million. The implied social cost for a given dollar of increased revenue is therefore 2.4. These results imply that even with high retailer market power, consumers are still deriving a large share of the benefits from this industry.

5.2 Is the Current Policy Maximizing Revenue?

In this section we examine whether the current excise tax is set at the revenue maximizing tax rate given that firms can respond to any tax change by strategically lowering their prices. Raising revenue for public use is cited as a primary justification for legalizing marijuana by every jurisdiction that has done so. The discussion that follows borrows from Miravete, Thurm, and Seim (2018a), who also consider the question of what tax rate maximizes revenue in the setting of excise taxes on alcohol in Pennsylvania. We differ from their approach in that we present results below for a model with multiple asymmetric retailers (instead of wholesalers in their case), each selling multiple products. This analysis also highlights the role of market power that the retailers have on the tax revenues. Note that the following analysis assumes that the processors do not have market power because, as we see in Section 3, there are more than 800 processors in the state and the extent of

²⁶These results are along the lines of what Conlon and Rao (2019) find in the liquor industry, in which consumers bear between 75% to 80% of the tax burden.

wholesale price variation is small. Hence, processors do not respond to price changes by retailers.

Single Product Monopoly In order to demonstrate how market power alters the excise tax design, we start from a simple set-up in which there is a single product monopoly retailer. The retailer's profit function is

$$\pi_r = (p^r - p^w)D((1 + \tau)p^r),$$

where p^r is the retail price and p^w is the wholesale price. Note that consumers pay $(1 + \tau)p^r$. The FOC of the retailer's optimization problem is

$$\frac{\partial \pi_r}{\partial p^r} = (p^r - p^w) \frac{\partial D((1 + \tau)p^r)}{\partial p^r} (1 + \tau) + D((1 + \tau)p^r) \quad (5.4)$$

and in equilibrium $\partial \pi_r / \partial p^r = 0$.

Applying the Implicit Function Theorem to equation 5.4, the tax pass-through rate can be written as

$$\frac{dp^r}{d\tau} = \frac{\kappa(p^*) - (2 - \frac{p^w}{p^r})}{p^*(2 - \kappa(p^*))}, \quad (5.5)$$

where $p^* = (1 + \tau)p^r$ and $\kappa(p)$ is the curvature of the demand curve, i.e., $\kappa(p) = \frac{D''(p)D(p)}{[D'(p)]^2}$. Moreover, the elasticity of the tax rate would be

$$\eta(\tau) = \frac{\partial p^r}{\partial \tau} \times \frac{\tau}{p^r} = -\frac{\tau}{1 + \tau} \times \frac{(1 - \frac{1}{\varepsilon(p^*)}) - \kappa(p^*)}{2 - \kappa(p^*)}. \quad (5.6)$$

Thus in the simple model, the degree to which taxes will be passed through to consumers in the form of higher prices depends on the elasticity of demand $\varepsilon(p)$ and the curvature of demand $\kappa(p)$. This latter measures how log-convex demand is. Intuitively, if demand is highly log-convex or curved, then when the tax rate goes up firms will respond by selling to a smaller but more inelastic population and will potentially raise prices by more than the amount of the tax increase.

Similarly, in this model the pass-through of wholesale price can be written as

$$\frac{dp^r}{dp^w} = \frac{1}{2 - \frac{DD''}{D'^2}} = \frac{1}{2 - \kappa(p^*)}, \quad (5.7)$$

which can be written as a function of the demand curvature, κ . Combining equations (5.6) and

(5.7), we can rewrite the tax elasticity as

$$\eta(\tau) = -\frac{\tau}{1+\tau} \times \left[\left(1 - \frac{1}{\varepsilon(p^*)}\right) - \left(2 - \frac{1}{\frac{dp^r}{dp^w}}\right) \right] \times \frac{dp^r}{dp^w}, \quad (5.8)$$

where $\varepsilon(p^*)$ is the demand elasticity of price evaluated at p^* . Hence, the elasticity of retail price with respect to tax depends on the elasticity of demand, ε and pass-through, $\frac{dp^r}{dp^w}$.

Now, we derive the revenue maximizing tax. Tax revenue is $R(\tau) = \tau p^r D((1+\tau)p^r)$ and the revenue maximizing tax satisfies

$$R'(\tau) = p^r D(p^*) \left[1 + \frac{\tau}{1+\tau} \varepsilon(p^*) + \eta(\tau)(1 + \varepsilon(p^*)) \right] = 0. \quad (5.9)$$

Hence, $R'(\tau) < 0$ if

$$1 + \frac{\tau}{1+\tau} \varepsilon(p^*) + \eta(\tau)(1 + \varepsilon(p^*)) < 0.$$

Note that the sign of $R'(\tau)$ is theoretically ambiguous and hence an empirical question. It depends on whether or not demand is sufficiently elastic relative to how much retailers will adjust their prices when the tax changes. Equations 5.6 and 5.9 show how calculating the revenue-maximizing tax rate can be made substantially more straightforward using empirically observed pass-through directly, rather than performing the calculation in equation 5.7 as the logit error imposes a particular restriction on the curvature of the demand curve.

It is useful at this point to compare how the government should set the tax differently under perfect competition and under imperfect competition. Under perfect competition, each retailer is a price taker and cannot affect the equilibrium price. In other words, $\eta(\tau) = 0$. Hence, the government increases the tax rate (i.e., $R'(\tau) > 0$) if and only if $1 + \frac{\tau}{1+\tau} \varepsilon(p^*) < 0$. The revenue-maximizing tax rate can then be set such that $\varepsilon(p^*) = -\frac{1+\tau}{\tau}$. This implies that for a 37% tax rate, as long as $\varepsilon((1+\tau)p^r) > -3.7$ the industry would be on the upward sloping portion of the Laffer curve. As shown by Anderson, de Palma, and Kreider (2001a), $\eta(\tau) > 0$ for a wide range of models. Hence, we can show that if the state can increase its tax revenue under imperfectly competitive market, then the state can also increase the tax revenue in the perfectly competitive market. Since we find a category-wide elasticity of -1.08 as shown in Table 8, the industry would clearly be on the upward sloping portion of the Laffer curve in the monopoly case.

Multi-product Oligopoly We extend the previous model to a more general case with multiple asymmetric retailers, each selling multiple products. With J retailers and K manufacturers transacting L products. Retailer g 's profit function is

$$\pi_g = \sum_{j \in J_g} (p_j^r - p_j^w) D_j((1 + \tau)p)$$

where j denotes product, J_g denotes the set of products that retailer i sells, p_j^r is retailer price of product j charged by retailer g , p_j^w is wholesale price of product j paid by retailer g , and p is a $J \times 1$ vector of retail prices $\{p_j^r\}$.

Under multi-product oligopoly, the state tax revenue is

$$R(\tau) = \tau \sum_j p_j^r D_j((1 + \tau)p^r).$$

The FOC of the tax-revenue maximization problem is

$$R'(\tau) = \sum_j p_j^r D_j \left[1 + \frac{\tau}{1 + \tau} \sum_k \varepsilon_{jk}(p^*) + \eta_j(\tau) + \sum_k \varepsilon_{jk}(p^*) \eta_k(\tau) \right], \quad (5.10)$$

where $\varepsilon_{jk}(p^*)$ is the demand elasticity with respect to retail price and η_j is the elasticity of the retail price with respect to the excise tax, i.e., $\eta_j(\tau) = \frac{\partial p_j}{\partial \tau} \frac{\tau}{p_j}$. In Appendix C we present a full derivation of results on $R'(\tau)$ in this general setting.

Like in the simple case, the elasticity of demand, $\varepsilon_{kj}(p^r)$, and the elasticity of price with respect to tax, $\eta_j(\tau)$ are directly informative on the sign of $R'(\tau)$ and evaluating this term is made substantially easier with estimated pass-through. This sign still depends on the curvature of demand but now also depends on consumer substitution within and across retailers and the relative margins of all the retailer's products.

If firms do not adjust prices and $\eta(\tau) = 0$, with $\tau = 0.37$, $R'(\tau) > 0$ if the aggregate elasticity of product j , $\sum_k \varepsilon_{jk} > -3.7$ for all j . In other words, if the market is perfectly competitive, the state is on the "right" side of the Laffer curve as long as the aggregate demand is sufficiently elastic. Using our demand estimates, we can calculate $R'(0.37)$ based on equation 5.9. We find that $R'(0.37)$ is significantly greater than 0. That is, our results indicate that the current excise tax is not too high to maximize tax revenue.

5.3 Counterfactual Policy Simulations

In this section we consider alternative regulatory arrangements including a state monopoly on retail sales. We also compute tax revenue under higher and lower tax rates as well as the effect be on total consumption and retailer profits. To evaluate these counterfactual policies, we need to impose a model of supply side competition between retailers. This will allow us to calculate how retailers will adjust prices in response to a tax or regulatory change. We incorporate estimated consumer demand and observed wholesale prices and assume that retailers set Nash-Bertrand prices.²⁷ This is a standard assumption in industrial organization, and is consistent with our conduct parameter estimate $\hat{\theta} = .89$, which implies each retailer behaves as if it is a local monopolist. Moreover, note that we show that the predicted pass-through based on the demand estimates and observed pass-through are close to each other. Although one typically uses *estimated* marginal costs in addition to estimated demand, we use the observed wholesale prices, which we set fixed in the counterfactual simulations as in the analysis in the previous section.

Table 11: Market Structure Counterfactuals (2017)

	Current Policy	Single-category Retailers	State Monopoly No Tax	State Monopoly With Tax	51.9% Margin
Avg. Pre-tax Price (\$)	15.63 (1.56)	11.3 (1.06)	18.28 (2.14)	17.8 (2.26)	8.10 (1.19)
Tax Revenue (millions of \$)	27.3 (7.96)	28.0 (7.77)	0	26.4 (7.56)	32.2 (11.47)
Usable Sales (kg)	3278 (321.90)	3947 (340.91)	3461 (361.22)	2508 (234.96)	6449 (575.50)
Retailer Revenue (millions of \$)	73.8 (21.53)	75.7 (21.01)	109.6 (33.17)	71.3 (20.44)	80.5 (30.99)
Retailer Variable Profits (millions of \$)	48.3 (13.22)	38.7 (9.99)	74.5 (20.37)	49.5 (13.71)	27.5 (10.59)
Consumer Surplus (millions of \$)	100.9 (25.58)	151.9 (39.03)	118.5 (28.21)	76.0 (18.51)	220.8 (51.89)

Note: The table shows monthly sales and Washington state tax revenue under the current policy as well as 4 counterfactual market structures: 1) retailers limited to selling single product categories to increase competition 2) a state monopoly with no 37% tax 3) a state monopoly with the 37% tax as well and 4) prices regulated to a uniform 51.9% margin above wholesale prices. Each value is calculated for 2017. Bootstrap standard errors are shown below in parentheses. They are calculated by drawing 100 times from the multivariate normal distribution of the parameter estimates for consumer demand and re-computing equilibrium outcomes at each set of parameters, then calculating the standard deviation.

²⁷We solve for the equilibrium set of Nash-Bertrand prices by iterating over the optimal markup for each firm in each market. We use a modified version of the fixed point problem suggested by Morrow and Skerlos (2010), where we modify the firms first order conditions to account for the excise tax on marijuana products. This tax creates a potentially large wedge between consumer prices and retailer profits and needs to be accounted for. We hold wholesale prices fixed in the counterfactuals in light of the competitive nature of the upstream market, described in the previous section.

Market Structure Counterfactuals In addition to tax policy, the state can regulate the market structure of the marijuana industry directly. An alternative policy available to Washington state would be to regulate the industry in the same way it had regulated liquor sales prior to 2012. The state had maintained a monopoly on retail sales of liquor and used a state-wide uniform markup of 51.9%. Other states, such as Pennsylvania, still regulate liquor in this way. In addition, some U.S. states have considered state monopolies on marijuana sales and 5 Canadian provinces are implementing government monopolies on marijuana retail. We test the counterfactual effects if the state were to switch to this policy for marijuana and show the results in column 5 of Table 11. These results use a markup of 51.9% which is the markup previously charged for liquor sales and currently used in Pennsylvania. At current wholesale prices, a markup of 51.9% translates to a margin of .36. This margin is substantially lower than the actual retail markups we observe in the data and so the average retail price is substantially lower in this counterfactual and total consumption increases substantially by 96% as well. We find that tax revenue would increase by 18%.

Because the observed retail markups are much higher than 51.9%, it may be more sensible to consider what a state monopoly would charge with no price regulation. In this case the state monopolist is the price-setter that would charge the profit maximizing price. While we assume the state monopolist would charge the profit-maximizing price, it is also possible the state-owned retailers would charge competitive prices, in which case outcomes would resemble column 1 but with retailer variable profits at least partially retained by the state. Results from this counterfactual are shown in column 3 and 4 of Table 11. We consider two possible policies, one in which the state has a monopoly on retail sales and maintains the current 37% tax and another where the tax is removed but the state earns the retail variable profits as revenue.²⁸

With the 37% tax still in place, we find the state retailer would increase prices, but only slightly. This follows from the fact that consumers do not appear to search actively and the private retailers already behave almost like local monopolists. Purely in terms of pricing then, consumers would be only mildly worse off. Total sales would fall slightly as would tax revenue. The state would earn substantially more in combined tax revenue and retail profits than by taxing private stores alone. The combined revenue has an upper bound of \$910 million per year, compared with \$327 million per year in revenue under the current policy.²⁹ If the state eliminated the tax as a source of revenue

²⁸We refer here to retail variable profits, calculated as retailer revenue minus cost of goods sold. Total profits after subtracting fixed costs including rent, labor costs, costs of federal income taxes, etc, will be significantly lower.

²⁹Annual revenue figures calculated here simply extrapolate the monthly figures from Table 11 for 12 months.

but kept all retail profits under the state monopoly, this would be slightly lower, at up to \$894 million in yearly profits. In other words, if the state’s goal is to raise revenue and maintain control over marijuana sales, monopolizing the retail industry directly would be much more lucrative than simply taxing retail sales at 37%.

By contrast, states may wish to reduce retailer market power if this results in lower prices, higher sales, and higher tax revenue. One way to do so would be to allow more retail entry or to restrict retailers to selling single product categories. Currently, as our demand estimates imply, retailers act almost as local monopolists and do not compete strongly on prices. Within a store there is significant competition between categories, however. If retailers were broken up they would no longer internalize this pricing externality. We test this counterfactual in Table 11 and find that prices would fall significantly, total sales would increase and total tax revenue go up slightly, increasing by roughly 3% to \$336 million per year.

Tax Policy Counterfactuals The previous section concluded that based on estimated price elasticity, it is highly likely that Washington state is still on the upward sloping region of the Laffer curve despite having the nation’s highest marijuana tax. To quantify the potential gains from increasing this tax rate further, we perform a set of counterfactual simulations considering increases in the tax rate from 37% to 40% and 50% as well as lowering it to 15%. For each tax rate, we allow firms to react to the tax change and re-solve for the Nash-Bertrand equilibrium in prices.

Results are presented in Table 12 for 2017. We show results both when firms respond strategically by changing prices and under an alternative where retailers lack market power and thus lack the ability to respond strategically. We find that firms would indeed respond to the tax change by decreasing pre-tax prices, and that increasing the tax rate to 40% would increase tax revenue by \$1.5 million per month, a 5.9% increase. Increasing the tax rate to 50% would increase tax revenue by \$5.9 million per month, a 22% increase. Comparing these results with the sufficient statistic results in the previous section, we find the social cost of raising the tax rate is nearly double the amount of revenue raised, with firms and consumers being harmed in roughly equal proportions. While the sufficient statistic approach predicts a social cost of taxation of 2.4 for every dollar raised, here we find a ratio of slightly under 2. The difference is because the sufficient statistic approach evaluates a unit tax and here we analyze a change in the ad valorem tax rate. As shown by Anderson, de Palma, and Kreider (2001b), when firms have market power an ad valorem tax is generally

Table 12: Counterfactual Tax Policy (2017)

	15%	37%	40%	50%
Monthly Tax Revenue (millions of \$)				
Strategic Price Reaction	13.5 (4.03)	27.3 (7.96)	28.8 (8.37)	33.2 (9.54)
Fixed Prices	13.7 (4.13)	27.3 (7.96)	28.6 (8.33)	32.4 (9.32)
Average Pre-Tax Price				
Strategic Price Reaction	17.2 (1.86)	15.6 (1.56)	14.8 (1.53)	14.0 (1.42)
Fixed Prices	14.5 (1.56)	14.5 (1.56)	14.5 (1.56)	14.5 (1.56)
Consumer Welfare				
Strategic Price Reaction	109.0 (24.58)	100.9 (25.58)	99.9 (27.53)	96.7 (24.58)
Fixed Prices	123.8 (29.09)	100.9 (25.58)	98.2 (24.92)	89.5 (22.88)
Retailer Variable Profit				
Strategic Price Reaction	61.9 (16.99)	48.3 (13.22)	46.8 (12.81)	42.3 (11.58)
Fixed Prices	70.4 (16.16)	48.3 (13.22)	46.0 (12.87)	39.2 (11.76)
Total Usable Sales (kg)				
Strategic Price Reaction	3370 (344.72)	3278 (321.90)	3266 (319.52)	3221 (312.71)
Fixed Prices	3770 (376.40)	3278 (321.90)	3214 (315.72)	3000 (296.77)

Note: The table shows monthly tax revenue and average prices under the current sales tax rate of 37% and the counterfactual rates of 15%, 40%, and 50% when firms are allowed to strategically respond to the tax increase by adjusting prices and when prices are fixed. Each value is calculated with data from 2017. Bootstrap standard errors are shown below in parentheses. They are calculated by drawing 100 times from the multivariate normal distribution of the parameter estimates for consumer demand and re-computing equilibrium outcomes at each set of parameters, then calculating the standard deviation.

welfare-superior.

We also find that increasing the tax rate to 50% would cause retail sales of usable marijuana to fall by approximately 60kg, or 2% of total sales.³⁰ If firms did not strategically reduce prices sales would instead fall by 278kg. Thus a naive policymaker not accounting for firm market power would predict a fall in consumption roughly 4.5 times larger than what would actually occur, in addition to underestimating the increase in revenue the tax change would cause.

These results suggest that Washington could significantly increase revenue by raising the tax rate, in part because retailers would respond to the tax by lowering their margins. We compare the expected increase in revenue when firms exercise their market power by strategically lowering pre-tax prices when the tax rate goes up. If retailers lack market power and do not adjust prices revenue will increase by roughly \$1 million fewer dollars per month, or a 17% smaller increase than if prices fully adjust. This illustrates that market power plays a significant role in how revenue will respond to a tax increase. If policymakers naively assume firms will not adjust prices in response to a tax change, their forecast of revenue will be significantly below the actual increase.

Finally, we evaluate how much revenue Washington would lose out on if it charged a 15% excise tax rate. As shown in Table 1, this is a common tax rate charged by many states, including California and Colorado. We estimate that under a 15% tax rate Washington would see monthly revenue of \$13.5 million, less than half of its current revenue. On an annual basis this would amount to \$165 million in foregone revenue in 2017. A simple extrapolation of this result to California, a state that taxes at 15%, implies that California is missing out on over \$800 million in annual revenue by under taxing marijuana relative to Washington's current 37% rate. This extrapolation assumes per capita marijuana demand is the same in the two states and that California is as successful as Washington at closing down black market retailers.³¹

Washington state's regulatory goals are to raise tax revenue and to restrain overall consumption. These results suggest the current regulatory regime is highly effective at reaching this goal given the high rate of pass-through and high retail margins. In addition, despite having the nation's highest tax rate, Washington is clearly on the upward sloping portion of the Laffer curve and could generate significantly higher revenue by increasing the tax rate. A higher tax rate would modestly

³⁰This can be taken as an upper bound on the increase in black market marijuana consumption following the tax increase under the worst case scenario where the entire decline in sales is explained by substitution to the black market. For reasons discussed previously in the paper we think this is unlikely.

³¹This also ignores revenue from license fees and the unit tax on production levied in California.

reduce marijuana consumption as well.

6 Conclusion

This paper studies the retail cannabis industry in the state of Washington, which was legalized in 2012 as the first state in the United States. Due to concerns over public health issues the state imposes tight regulation over marijuana consumption similarly to other sin-product markets such as alcohol and tobacco. In particular, state tax on retail sales is 37% in Washington, which is higher than any other states that have legalized recreational marijuana sales, and tight retail license cap limits fierce competition among retailers. The main purposes of the regulatory framework are increasing tax revenue from marijuana sales and controlling marijuana consumption at the same time.

We use detailed transaction-level wholesale price and retail price data to investigate the incidence of these taxes and whether the state is overtaxing the product and reducing revenue, as well as the regulatory design of the retail market under market power. It is important to examine the effect of market power because the retail license cap limits competition and allows retailers to sustain high margins. Moreover, most studies of taxation in public finance consider perfectly competitive markets. Hence, the literature studying the role of market power in taxation is still very scarce.

Our analysis proceeds in four steps. First, we estimate consumer demand, which we model in the horizontally-differentiated product framework following Berry, Levinsohn, and Pakes (1995). Our demand estimates imply that consumer cannabis demand is relatively elastic and retailers have significant market power partially due to the entry restriction that the state imposes. Second, we estimate conduct parameters by comparing observed margins to the margins implied by the highly elastic demand. We use these as a sufficient statistic for competition when estimating tax incidence. Third, we estimate cost pass-through, which is a key input for calculating tax incidence following the method proposed by Fabinger and Weyl (2013). Since our data contain detailed information on wholesale prices, neither of these require estimation of marginal costs. We find that costs are more than fully passed through from retailers to consumers. Lastly, combining three pieces together, we provide extensive policy analysis. In particular, we calculate the tax incidence and the social cost of tax. Moreover, we conduct a series of counterfactual simulations to highlight the role of competition in designing sales taxes. Our results indicate that despite having the nation's highest

tax rate, Washington still has significant scope to increase revenues with a higher tax rate. That is, they are still on the left side of the Laffer curve. We also find significant social costs of taxation, more than 2 dollars are lost to consumers and producers for every dollar of tax revenue generated. Lastly, we find that the state can increase the degree of competition by, for example, increasing the license cap in order to increase tax revenue.

There are some interesting issues that may worth studying in the future. For example, we abstract away from dynamics in both consumer and firm behavior. Similarly to other sin products, addiction to marijuana is an important concern for the state, but our current demand model does not allow explicit inter-temporal linkage through addiction. Also, retailers need to learn consumer demand and competitor behavior in a newly created market as in the legalized marijuana market. Studying both demand- and supply-side dynamics would be a fruitful topic for the future research.

References

- Adda, J., B. McConnell, and I. Rasul (2014): “Crime and the Depenalization of Cannabis Possession: Evidence from a Policing Experiment,” *Journal of Political Economy*, 122(5), 1130–1202.
- Agarwal, S., S. Chomsisengphet, N. Mahoney, and J. Stroebel (2014): “A Simple Framework for Estimating Consumer Benefits from Regulating Hidden Fees,” *The Journal of Legal Studies*, 43(S2), S239–S252.
- Aguirregabiria, V., D. Ershov, and J. Suzuki (2015): “Estimating the Effects of Deregulation in the Ontario Alcohol Retail Market,” *Working Paper*.
- Allcott, H., B. Lockwood, and D. Taubinsky (2019): “Regressive Sin Taxes, with an Application to the Optimal Soda Tax,” *Quarterly Journal of Economics*, pp. 1557–1626.
- Anderson, S., A. de Palma, and B. Kreider (2001a): “Tax Incidence in Differentiated Product Oligopoly,” *Journal of Public Economics*, 81, 173–192.
- (2001b): “The Efficiency of Indirect Taxes under Imperfect Competition,” *Journal of Public Economics*, 81, 231–251.
- Atkin, D., and D. Donaldson (2015): “Who’s Getting Globalized? The Size and Implications of Intra-national Trade Costs,” *Working Paper*.
- Atkinson, A., and J. Stiglitz (2015): *Lectures on Public Economics*. Princeton University Press.
- Auerbach, A. (1985): “The Theory of Excess Burden and Optimal Taxation,” *Handbook of Public Economics*, 1.
- Becker, G., M. Grossman, and K. Murphy (1991): “Rational Addiction and the Effect of Price on Consumption,” *American Economic Review: Papers and Proceedings*, 81(2), 237–241.
- Berry, S. (1994): “Estimating Discrete Choice Models of Product Differentiation,” *RAND Journal of Economics*, 25, 242–262.
- Berry, S., and P. Haile (2014): “Identification in Differentiated Products Markets Using Market Level Data,” *Econometrica*, 82, 1749–1797.
- Berry, S., J. Levinsohn, and A. Pakes (1995): “Automobile Prices in Market Equilibrium,” *Econometrica*, 63, 841–890.
- Bollinger, B., and S. Sexton (2018): “Local Excise Taxes, Sticky Prices, and Spillovers: Evidence from Berkeley’s Soda Tax,” *mimeo*.
- Cawley, J., and D. Frisvold (2017): “The Incidence of Taxes on Sugar-Sweetened Beverages: The Case of Berkeley, California,” *Journal of Policy Analysis and Management*, 36(2).
- Chamberlain, G. (1987): “Asymptotic Efficiency in Estimation with Conditional Moment Restriction,” *Journal of Econometrics*, 34, 305–334.
- Conlon, C., and J. Gortmaker (2019): “Best Practices for Differentiated Products Demand Estimation with pyblp,” *Working Paper*.

- Conlon, C., and N. Rao (2019): “Discrete Prices and the Incidence and Efficiency of Excise Taxes,” *Working Paper*.
- Davis, C., M. Hill, and R. Phillips (2019): “Taxing Cannabis,” *Institute on Taxation and Economic Policy*.
- Dube, J. P., and S. Gupta (2008): “Cross-Brand Pass-Through in Supermarket Pricing,” *Marketing Science*, 27(3), 324–333.
- Fabinger, M., and G. Weyl (2013): “Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition,” *Journal of Political Economy*, 121.
- Fabra, N., and M. Reguant (2014): “Pass-Through of Emissions Costs in Electricity Markets,” *American Economic Review*, 104(9), 2872–99.
- Gavrilova, E., T. Kamada, and F. Zoutman (forthcoming): “Is Legal Pot Crippling Mexican Drug Trafficking Organisations? The Effect of Medical Marijuana Laws on US Crime,” *The Economic Journal*.
- Gordon, B., and B. Sun (2015): “A Dynamic Model of Rational Addition: Evaluating Cigarette Taxes,” *Marketing Science*, 34(3), 309–472.
- Grigolon, L., and F. Verboven (2014): “Nested Logit or Random Coefficients Logit? A Comparison of Alternative Discrete Choice Models of Product Differentiation,” *Review of Economics and Statistics*, 66, 916–935.
- Hansen, B., K. Miller, and C. Weber (2018): “The Taxation of Recreational Marijuana: Evidence from Washington State,” *Working Paper*.
- Hansen, B., K. Miller, and C. Weber (2020): “Early Evidence on Recreational Marijuana Legalization and Traffic Fatalities,” *Economic Inquiry*, 58(2), 547–568.
- Hanssens, D. (2015): *Empirical Generalizations About Marketing Impact*. Marketing Science Institute.
- Hao, Z., and B. Cowan (2017): “The Cross-border Spillover Effects of recreational marijuana legalization,” *mimeo*.
- Hong, G. H., and N. Li (2017): “Market Structure and Cost Pass-Through in Retail,” *The Review of Economics and Statistics*, 99(1), 151–166.
- Hottman, C. (2018): “Retail Markups, Misallocation, and Store Variety across U.S. Cities,” *Working Paper*.
- Jacobi, L., and M. Sovinsky (2016): “Marijuana on Main Street? Estimating Demand in Markets with Limited Access,” *American Economic Review*, 106.
- Khan, R., K. Misra, and V. Singh (2016): “Will a Fat Tax Work?,” *Marketing Science*, 35(1), 10–26.
- Levy, D., D. Lee, H. A. Chen, R. J. Kauffman, and M. Bergen (2011): “Price Points and Price Rigidity,” *The Review of Economics and Statistics*, 93(4), 1417–1431.

- McShane, B. B., C. Chen, E. T. Anderson, and D. I. Simester (2016): “Decision Stages and Asymmetries in Regular Retail Price Pass-Through,” *Marketing Science*, 35(4), 619–639.
- Miller, K., B. Hansen, and C. Weber (2018): “Federalism, Partial Prohibition, and Cross Border Sales: Evidence from Recreational Marijuana,” *mimeo*.
- Miller, K., and B. Seo (2019): “Tax Revenues When Substances Substitute: Marijuana, Alcohol, and Tobacco,” *Working paper*.
- Miller, N., M. Osborne, and G. Sheu (2017): “Pass-through in a concentrated industry: empirical evidence and regulatory implications,” *The RAND Journal of Economics*, 48(1), 69–93.
- Miller, N., and M. Weinberg (2017): “Understanding the Price Effects of the MillerCoors Joint Venture,” *Econometrica*.
- Miravete, E., J. Thurk, and K. Seim (2018a): “Market Power and the Laffer Curve,” *Econometrica*, 86(5), 1651–1687.
- (2018b): “One Markup to Rule them All: Taxation by Liquor Pricing Regulation,” *Working Paper*.
- Morrow, W. R., and S. Skerlos (2010): “On the Existence of Bertrand-Nash Equilibrium Prices Under Logit Demand,” *CoRR*, arXiv:1012.5832.
- Nakamura, E., and D. Zerom (2010): “Accounting for Incomplete Pass-Through,” *The Review of Economic Studies*, 77(3), 1192–1230.
- Reynaert, M., and F. Verboven (2014): “Improving the performance of random coefficients demand models: The role of optimal instruments,” *Journal of Econometrics*, 179, 83–98.
- Seiler, S., A. Tuchman, and S. Yao (2018): “The Impact of Soda Taxation,” *mimeo*.
- Thomas, D. (2018): “License Quotas and the Inefficient Regulation of Sin Goods: Evidence from the Washington Recreational Marijuana Market,” *mimeo*.
- Thomassen, O., H. Smith, S. Seiler, and P. Schiraldi (2017): “Multi-Category Competition and Market Power: A Model of Supermarket Pricing,” *American Economic Review*, 107(8), 2308–2351.
- Varadhan, R., and C. Roland (2008): “Simple and globally convergent methods for accelerating the convergence of any EM algorithm,” *Scandinavian Journal of Statistics*.
- Waldfogel, J., and K. Seim (2013): “Public Monopoly and Economic Efficiency: Evidence from the Pennsylvania Liquor Control Boards Entry Decisions,” *American Economic Review*, 103.

Appendix A Data Cleaning and Price Analysis

In this appendix, we describe how the dataset was cleaned and document additional details on retail and wholesale prices discussed in Section 3.

A.1 Application Data

The list of businesses that have applied to licenses is available at the Washington State Liquor and Cannabis Board website.³² This list of licenses is not cumulative as we noticed that some firms that did not obtain licenses are dropped from the file through time. To recover the history of all applications we use the Internet 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.

A.2 Transaction Data

We have four distinct data sets that are put together to form the final 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 in cases where firms own multiple licenses.
- 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.
- Lab results samples: contains information regarding samples. It links the sample id to parentid.

³²<https://data.lcb.wa.gov>

- Lab results potency: contains potency analysis (THC and CBD content) by sample id, but not the inventory id used the match transfers and transactions.

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 first two datasets above.

A.3 Taxes and 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. 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).³³

We also 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.

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.

³³Legal purchase limits are one ounce for usable, 16 ounces for solid, 72 ounces for liquid, and 7 grams for concentrates.

Table A.1: Pass-through Estimates: First Difference with Weekly Prices

	(1)	(2)	(3)	(4)
d_whole	1.528*** (0.00192)	1.515*** (0.00212)	1.563*** (0.00245)	1.515*** (0.00283)
d_thc	0.00779*** (0.000522)	0.00848*** (0.000577)	0.00869*** (0.000689)	0.00992*** (0.000871)
d_whole_comp		0.0112 (0.0102)	0.0108 (0.0135)	0.0898*** (0.0226)
Constant	-0.0345*** (0.00266)	-0.0316*** (0.00294)	-0.0418*** (0.00379)	0.0180** (0.00616)
Year, month, week dummy	x			
Year x Week dummy		x	x	x
Exclude d_retail = 0			x	
Exclude d_whole = 0				x
Observations	4506187	3673076	2841230	1485297
R-squared	0.13	0.12	0.13	0.16

Note: The table reports pass-through estimates with weekly average prices based on the first-difference linear panel model. Standard errors are clustered at the city level. Model (3) uses the observations only when retailers change prices. Model (4) uses the observations excluding the case where processors do not change prices. Significance levels: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Appendix B Additional Results on Prices and Pass-through

B.1 Pass-through with Weekly Prices

The pass-through estimates in the main text use the retail prices and wholesale prices aggregated at the monthly level. To see how this aggregation affect our pass-through estimates, we now estimate equation (4.2) with prices aggregated at the weekly level.

Table A.1 reports the estimation results and we find that the coefficient of Δw_{gjt} is around 1.5 for all specifications, which is close to the results in Table 6. Model (1) controls for the change in THC and year, month, and week dummies. Model (2) includes the change in the competitors' wholesale prices, and the interaction of year and week dummies. Model (3) excludes the observations

Figure A.1: Retail Tax-inclusive List Price Distribution by Category

when retailers do not change the price, and Model (4) excludes the observations when processors do not change the price. Thus, our pass-through estimates are robust to the level of aggregation.

B.2 Discrete Prices

B.2.1 Price Distribution

Discrete price changes have been documented in some retail industries such as liquor (Conlon and Rao (2019)) and grocery stores (Levy, Lee, Chen, Kauffman, and Bergen (2011)). In this section, we document *daily* price changes in the tax-inclusive posted prices at the retail stores of marijuana and examine the implications of the discrete prices for our pass-through estimates.³⁴

Figure A.1 plots the distribution of the tax-inclusive retail posted price by product category. Prices are daily prices instead of monthly average prices as we used in the demand estimation and the pass-through regression. As the figure shows, a large fraction of the retail posted prices are ending at 0. Table A.2 further confirms that 83.15% of transactions involve the retail price ending at zero.

Similarly, Figure A.2 shows the distribution of wholesale prices. The distribution of wholesale

³⁴For the monthly aggregated prices, we do not see any discrete price distribution.

Table A.2: Prices Ending at 0

	Tax Inclusive Retail Price	Wholesale Price
End at 0	83.15%	44.57%

Figure A.2: Wholesale Price Distribution by Category

prices is lower than the distribution of retail prices as we find in Section 3. Moreover, we find some mass points across categories but wholesale prices end with 0 with less frequency than retail prices. About 44.57% of wholesale prices end at 0.

B.2.2 Price Changes

Figure A.3 plots the histogram of the change in tax-inclusive retail posted prices. We find substantial price rigidities for retail prices. We find that retailers do not change prices in about 51% of the cases.

Figure A.4 shows the histogram of price changes in tax-inclusive retail prices conditional on a price change from the previous day. Conditional on a price change, retail prices tend to change by an \$1 increment. In fact, Table A.3 shows that about 60% of price changes are \$1 increments, which is smaller than what Conlon and Rao (2019) find. Note again that our price data is *daily*

Figure A.3: Retail Tax-inclusive List Price Changes by Category

price changes, whereas Conlon and Rao (2019) study *monthly* price changes.

Table A.3: Price Changes with \$1

	Change in Tax Inclusive Retail Price	Change in Wholesale Price
\$1 increment	60.75%	23.60%

Note: The numbers are based on the observations conditional on prices being changed.

B.2.3 Pass-through

Discrete price changes imply that estimated pass-through rates could be overestimated as suggested by Conlon and Rao (2019). To examine how discrete price changes affect pass-through estimates, we follow Conlon and Rao (2019) and estimate the following ordered logistic regression model.

$$\begin{aligned} \Delta p_{ijt} &= k \text{ if } \Delta p_{ijt}^* \in [\alpha_k, \alpha_{k+1}] \\ \Delta p_{ijt}^* &= f(\Delta w_{ijt}, \Delta \bar{w}_{-it}, \Delta x_{ijt}) + \mu_t + \varepsilon_{ijt}, \end{aligned} \tag{B.1}$$

Figure A.4: Retail Tax-inclusive List Price Changes by Category (Excluding no price changes)

where $f(\cdot)$ is a flexible function of $\Delta w_{ijt}, \Delta \bar{w}_{-it}, \Delta x_{ijt}$. We also include days since the last price change and whether a price increases or not, interacted with $(\Delta w_{ijt}, \Delta \bar{w}_{-it})$. Using the estimates,³⁵ we calculate the pass-through rate by predicting a change in the retail price following a change in wholesale price for each observation. In particular, we consider $\Delta \in \{\$0.01, \$0.1, \$0.5, \$1, \$2\}$.

Figure A.5 plots the predicted mean pass-through for different amount of price changes ranging from \$0.01 to \$2. We find that the mean pass-through slightly increases as the price change is bigger. The pattern is consistent with the finding by Conlon and Rao (2019).

Appendix C The Laffer Curve under Multi-product Oligopoly

In this appendix, we extend the results presented in section 5.2 and derive results on the relationship between the tax rate and total revenue for the general case with multiple asymmetric retailers, each selling multiple products.

³⁵We do not report estimated coefficients as the function is highly nonlinear.

Figure A.5: Predicted Pass-through by Different Price Changes

With G retailers and K manufacturers transacting J products. Retailer g 's profit function is

$$\pi_i = \sum_{j \in J_g} (p_j^r - p_j^w) D_j((1 + \tau)p)$$

where j denotes product, J_g denotes the set of products that retailer g sells, p_j^r is retailer price of product j charged by retailer g , p_j^w is wholesale price of product j paid by retailer g , and p is a $J \times 1$ vector of retail prices $\{p_j^r\}$. Note that $(1 + \tau)p_j^r$ is the retailer price that consumers actually pay.

The FOC of retailer j 's profit maximization problem is

$$D_j + \sum_{j' \in J_g} \left[(p_{j'}^r - p_{j'}^w)(1 + \tau) \frac{\partial D_{j'}}{\partial p_j^r} \right] = 0. \quad (\text{C.1})$$

As we have shown in Section 2, there are a large number of processors relative to the number of retailers that are capped by the regulation. Hence, we assume that the manufacturers do not have any market power and charge their marginal cost to retailers. This implies that wholesale prices do not respond to a change in the retail price. Given this assumption, applying the Implicit Function Theorem to equation (C.1) gives the (own) pass-through rate of wholesale prices to retail prices as follows:

$$\frac{dp_j^r}{dp_j^w} = \frac{\frac{\partial D_j}{\partial p_j^r}}{2\frac{\partial D_j}{\partial p_j^r} + \sum_{j' \in J_g} \left[(p_{j'}^r - p_{j'}^w)(1 + \tau) \frac{\partial^2 D_{j'}}{\partial p_j^{r^2}} \right]}. \quad (\text{C.2})$$

Similarly, the pass-through rate of excise tax to retail prices can be written as

$$\frac{dp_j^r}{d\tau} = \frac{\sum_k \frac{\partial D_j}{\partial p_k^r} p_k^r + \sum_{j'} \left[(p_{j'}^r - p_{j'}^w) \left((1 + \tau) \sum_k \frac{\partial \left(\frac{\partial D_{j'}}{\partial p_k^r} \right)}{\partial p_k^r} p_k^r + \frac{\partial D_{j'}}{\partial p_j^r} \right) \right]}{2\frac{\partial D_j}{\partial p_j^r} + \sum_{j' \in J_g} \left[(p_{j'}^r - p_{j'}^w)(1 + \tau) \frac{\partial^2 D_{j'}}{\partial p_j^{r^2}} \right]} \quad (\text{C.3})$$

Combining equation C.2 and equation C.3, we obtain

$$\frac{dp_j^r}{d\tau} = \sum_k \frac{\partial D_j}{\partial p_k^r} p_k^r + \sum_{j'} \left[(p_{j'}^r - p_{j'}^w) \left((1 + \tau) \sum_k \frac{\partial \left(\frac{\partial D_{j'}}{\partial p_k^r} \right)}{\partial p_k^r} p_k^r + \frac{\partial D_{j'}}{\partial p_j^r} \right) \right] \times \frac{\frac{dp_j^r}{d\tau}}{\frac{\partial D_j}{\partial p_j^r}} \quad (\text{C.4})$$

Hence, the pass-through of tax depends on the wholesale pass-through, demand elasticity and the curvature of the demand. Compared to the single-product monopoly case, one needs the information about demand curvature to calculate the tax pass-through.

Now, consider the tax revenue for the state of Washington from the sales of cannabis is

$$R(\tau) = \tau \sum_j p_j^r D_j ((1 + \tau) p^r)$$

The FOC of the tax-revenue maximization problem is

$$\begin{aligned} R'(\tau) &= \sum_j p_j^r D_j ((1 + \tau) p^r) + \tau \sum_j p_j^r \sum_k \frac{\partial D_j}{\partial p_k^r} p_k^r \\ &\quad + \tau(1 + \tau) \sum_j p_j^r \sum_k \frac{\partial D_j}{\partial p_k^r} \frac{dp_k^r}{d\tau} + \tau \sum_j \frac{dp_j^r}{d\tau} D_j \\ &= \sum_j p_j^r D_j \left[1 + \frac{\tau}{1 + \tau} \sum_k \varepsilon_{jk}(p^*) + \eta_j(\tau) + \sum_k \varepsilon_{jk}(p^*) \eta_k(\tau) \right], \end{aligned} \quad (\text{C.5})$$

where $\varepsilon_{jk}(p^*)$ is the demand elasticity with respect to retail price and η_j is the elasticity of the retail price with respect to the excise tax, i.e., $\eta_j(\tau) = \frac{\partial p_j}{\partial \tau} \frac{\tau}{p_j}$.

The optimal excise tax satisfies $R'(\tau) = 0$. We evaluate $R'(\tau)$ locally in the area around the

current tax rate to determine which side of the Laffer curve current policy resides. We do so by evaluating equation C.5 using its empirical counterparts estimated in the previous section. A key part of determining $R'(\tau)$ is the tax elasticity $\eta(\tau)$, which describes how retailers will adjust their prices in response to a tax change. Under the perfect competition, again, $\eta_j(\tau)=0$. Hence, $R'(\tau)<0$ if and only if $\sum_j p_j^r D_j \left(1 + \frac{\tau}{1+\tau} \sum_k \varepsilon_{jk}(p^*)\right) < 0$. Since $p_j^r D_j((1+\tau)p) > 0$, the sign of $R'(\tau)$ depends on the sign of $1 + \frac{\tau}{1+\tau} \sum_k \varepsilon_{jk}(p^*)$. Given that $\tau = 0.37$, $R'(\tau) > 0$ if the aggregate elasticity of product j , $\sum_k \varepsilon_{jk} > -3.7$ for all j . In other words, if the market is perfectly competitive, the state is on the “right” side of the Laffer curve as long as the demand is sufficiently elastic. Using only demand estimates, we can calculate $R'(0.37)$ based on equation C.5. We find that $R'(0.37)$ is greater than 0. That is, our results indicate that the current excise tax is not too high to maximize tax revenue.