Taxation and Market Power in the Legal Marijuana Industry

Brett Hollenbeck
Anderson School of Management
UCLA [*]

Kosuke Uetake
School of Management
Yale University[†]

This draft: October 2018[‡]

Abstract

In 2012 the state of Washington created a legal framework for production and retail sales of marijuana. Nine 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 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 examine 3 main research questions. First, how effective is Washington's excise tax at raising revenue? With the nation's highest tax rate on marijuana, is Washington maximizing revenue or potentially overtaxing, leading to reduced legal sales and lower tax revenue. Second, what is the incidence of taxes in this industry? Finally, most states have restricted entry, resulting in firms with substantial market power. What is the role of imperfect competition in studying these basic questions on tax policy? We combine structural methods and a reduced form sufficient statistic approach to show a number of results. First, Washington's 37% excise tax is still on the upward sloping portion of the Laffer curve and state revenue could be substantially higher with a higher tax rate. The amount of revenue generated by a tax increase is significantly larger due to retailer market power than it would be under perfect competition. In addition, these taxes are primarily borne by consumers and not by firms, and there is a large social cost associated with each dollar raised.

[*]brett.hollenbeck@anderson.ucla.edu

[†]kosuke.uetake@yale.edu

[‡]The authors thank the Morrison Center for Marketing Research for generous funding. We wish to thank Renato Giroldo for excellent research assistance. We also are grateful for comments from Sylvia Hristekeva, 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.

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

1 Introduction

By November 2018, ten 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 25% 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 2017, this industry accounted for $8.5 billion in sales in the U.S., a figure which is expected to grow to $57 billion in annual sales in the next decade, making it comparable to or larger than other "sin" products such as liquor or wine.[1]

Similarly to alcohol, states have chosen to tightly regulate this industry due to concerns over public health issues related to marijuana 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 created significant uncertainty among 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 6.5% in Massachusetts to 37% in Washington, nearly 5 times higher, illustrating this uncertainty. The stakes of this decision are large, as the consequence of a difference of this size for a large state

[1] Wine sales in the U.S. totaled $41 billion in 2017, liquor sales totaled $25 billion, and tobacco sales totaled $121 billion. Data on current sales and forecast for future marijuana sales growth come from Arcview Market Research and BDS Analytics.

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

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 black market consumption? This is effectively a Laffer Curve criterion and is widely cited by U.S. states and Canadian provinces as the primary reason to keep tax rates low.[3] Washington policymakers consider this an open debate, with legislation introduced in 2016 that would lower the tax rate and suggesting this would increase revenue.[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 enhance 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 but also shed light on this question.

Third, almost all 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 more firm market power. We therefore incorporate the role 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]

We are aided by an exceptionally rich and comprehensive new source of data. Washington

[3]The 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).

[4]HB 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."

[5]As 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."

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. 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 tax and regulatory policies, we impose a model of supply-side competition and verify that it replicates observed margins and cost pass-through.

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 $16.1 per gram and an average wholesale price of $6.9 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 (2001) show that the degree of monopoly power has a significant effect on the extent to which taxes will be passed through to consumers. In Washington, an increase in the tax rate might cause retailers to lower their margins, 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 marijuana 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 $2.5-2.9$. These are 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 -1.4, 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 4 that this elasticity implies the industry is still on the upward sloping region of the demand 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

[6] See 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.

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

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. We use these results on pass-through to measure the incidence of taxes in this new industry as well as the social cost of taxation. 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 27% falling on producers and the remaining 73% by consumers. These taxes can 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, we evaluate counterfactual tax rates to see how much additional revenue the state is forgoing with a 37% tax. Despite having the nation's highest tax rate for marijuana products at 37%, revenue could still be substantially higher. We show that increasing the tax rate from 37% to 40% would increase revenue by approximately $17 million per year, from $304 million to $321 million, and raising the tax rate to 50% would increase revenue by $66 million per year. On the other hand, if Washington set taxes at 15% like many other large states, annual revenue would be lower by $157 million, or roughly 50%.[9] Retailer

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

[9] 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 undertaxing marijuana relative to Washington's current 37% rate.

market power plays a significant role in this result. We compare the change in revenue when retailers strategically adjust prices following a tax increase to those where retailers act as price-takers and do not respond and find that the change in revenue is 30% higher due to retailer market power.

We next 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 $484 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.

This paper is related to several literatures. The first is the recent empirical literature on sin products that has focused on alcohol taxation and regulation and has used 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 (2016), 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 and market power resides primarily with retailers. Because there is no regulation of upstream margins, we can also estimate retail pass-through. We show how this contributes to identification of the likely effects of a tax change. Whereas they find Pennsylvania is on the wrong side of the Laffer Curve, we find that Washington is still on the upward sloping region.

[10]At 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 been 3.6% even without a tax increase.

Other 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.

Third, this paper relates to the new and growing literature on legal and illegal cannabis industries (see, e.g., Jacobi and Sovinsky (2016); Adda, McConnell, and Rasul (2014); Dragone, Prarolo, Vanin, and Zanella (2017)). Jacobi and Sovinsky (2016) use the 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. Adda, McConnell, and Rasul (2014) 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. Hansen, Miller, and Weber (2018) study the effects of the change in tax structure in Washington in 2015 and present results on the effects of taxes on vertical integration incentives and the short-term effects of the change on prices. (Thomas 2018) studies the welfare implications of license quota and find that allowing free entry raises the state's tax revenue relative to the current quota system. Hao and Cowan (2017) studies the spillover effects of recreational marijuana legalization (RML) 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.

This paper also 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 (2016)) 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.

2 Data and Industry Background

2.1 Regulation and Taxation

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.

[11]The 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. The stark contrast between how vertical integration is treated under these two regulatory regimes highlights the large degree of uncertainty policymakers have regarding how this new market should be best regulated.

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 beginning when a seed is planted and following it to the final retail sale.

The data contain 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. 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 limit on the number of processing licenses.

Table 1: Marijuana Excise Tax by States

State	Sales Tax	Annual Revenue
AK	$50/oz	
CA	15%	
CO	15%	$205 million
MA	6.25%	
ME	10%	
NV	15% wholesale, 10% retail	
OR	17%	$55 million
WA	37%	$281 million

[1] Other: CA $9.25/oz flowers & $2.75/oz leaves.
[2] Some localities also impose their own excise taxes.

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.[12] 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.[13]

Table 1 reports the sales tax for 8 states that have already started the legalized cannabis industry. As the table shows, the sales tax rate varies significantly across states, ranging from 6.25% in MA to 37% in WA. Washington charges the highest sales taxes on marijuana by a large margin.[14]

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

[13]Because 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.

[14]Washington 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

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

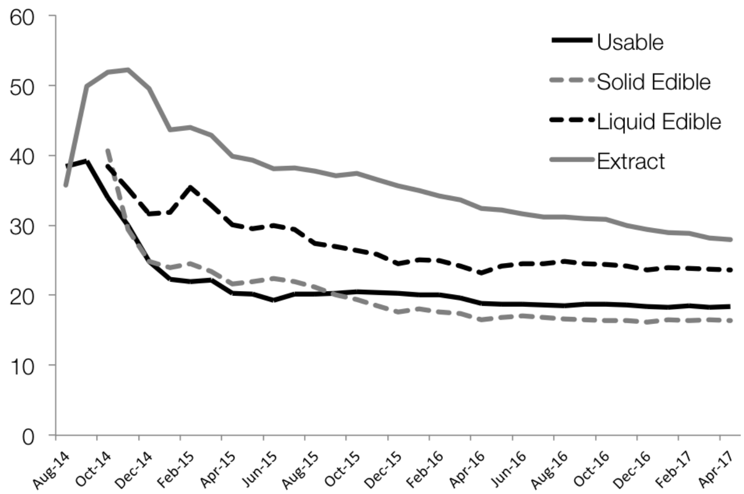

2.2 Descriptive Results

In this section we describe the key features of the data that motivate our empirical analysis.

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

highest tax on cigarettes at $3.025 per packof 20 cigarettes.

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

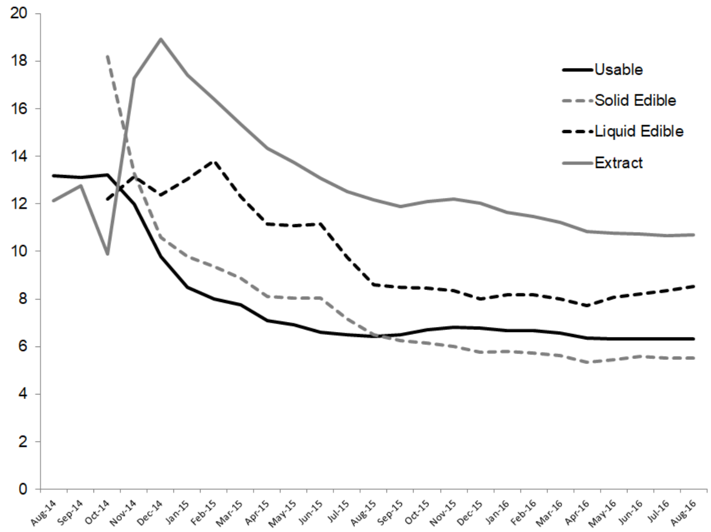

vapor. These account for 96% of sales, with the remaining 4% consisting of a large number of niche products which will largely be excluded from analysis. Within each of these product categories, there is some remaining product heterogeneity which is mostly unobserved.

Figure 1 plots the average (tax inclusive) retail price over 4 years in our data. Generally, retail prices decrease over time for all categories. The figure shows that in 2017, the average retail price across all products was $15.21 per gram excluding taxes, where 1 gram is a standard product unit. This corresponds to a mean price of $20.83 including taxes. 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 ($7.48 per gram) 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 $6.80 out of a total retail price of $13.49, yielding an average margin of .50 for usable products and .54 for all products. Aggregating at the level of product type, retailer margins ranged from .33 to .67

Table 2: Price Summary Statistics (2017)

		Retail Price		Wholesale Price	
	Total Sales (grams)	Mean	Std Dev	Mean	Std Dev
All Combined Products	5,052.9	15.2	5.0	7.5	2.5
Usable Marijuana	3,585.7	13.49	3.8	6.7	2.1
Solid Edible	400.2	12.2	2.6	5.9	1.3
Liquid Edible	101.2	17.5	3.8	8.6	1.7
Extract	726.3	20.8	3.7	10.4	1.5
Other Products	239.0	11.2	3.4	5.4	1.8

Note: This table presents total sales and average prices for each product type during the year 2017. 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

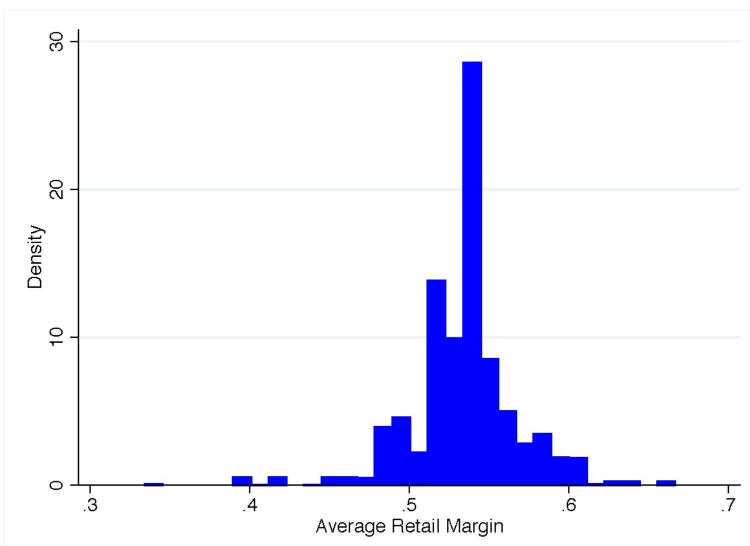

Figure 4: Average Monthly Sales by Type

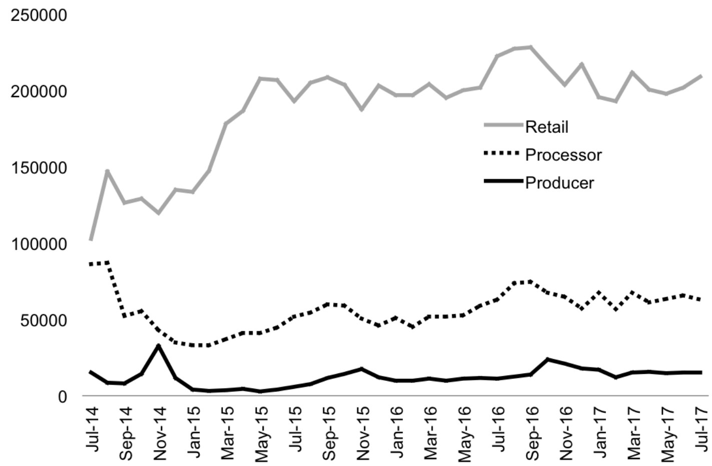

with most retailers setting margins between .5 and .6. These margins are substantially higher than typical margins in retail settings. The median margin for U.S. grocery products has been estimated at roughly .3 (Hottman (2018)) with higher estimates of .45 in the U.K. (Thomassen, Smith, Seiler, and Schiraldi (2017)) and with an upper bound of .52.

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. The industry has shown rapid then steady growth, with retailers averaging slightly more than $200,000 in monthly sales 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 July 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

Table 3: Firm Revenue in 2017

		Mean	5th Pctile	95th Pctile	Std	Max
Number of Firms (2017)						
Retailers	385					
Processors	642					
Producers	388					
Monthly Revenue (2017)						
Retailers		$202,354.8	$21,947.5	$573,506.4	185,117.8	$1,394,183.0
Processors		$63,377.6	$2,100.0	$247,165.2	149,459.5	$2,181,563.1
Producers		$14,921.3	$560	$69,038.4	22,975.3	$174,856.9

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.

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

Note: The table shows summary statistics of the transactions between retailers and processors.

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 second row shows the share of sales from each processor per retailer. The average market share is 8% and the median market share is 1.5%. Hence, retailers have a lot of transaction partners and are not dependent on any particular processor. These facts indicate that different processors are close substitutes from the retailer's perspective.

3 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.[15] 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 goals have nevertheless chosen very different tax rates and regulatory regimes.[16] 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 (2001) 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 (2016), 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

[15] Washington 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%.

[16] For 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.

of these in succession.

3.1 Demand Estimation

In this section we describe the method used to estimate consumer elasticity of demand in this industry. Measuring consumer price elasticity is necessary to understand how consumption and tax revenue would change under counterfactual taxes as well as the incidence of the current taxes. We follow the large literature on using market-share data to estimate demand as a function of product characteristics beginning with Berry (1994), Berry, Levinsohn, and Pakes (1995) (BLP), and Nevo (2001).[17]

We proceed with a model of random coefficient nested logit (RCNL) demand in order to produce robust own and cross-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 and to produce realistic own and cross-price elasticities. Market is defined at the city level and product is defined at the retailer-category level. Following the discrete choice demand literature, we model demand over $j \in \mathscr{J}$ products in each market in time period t for a set of consumers defined by i. Each consumer has utility which is modeled as

$$u_{ijt} = x_{jt}\beta + \alpha_i^* p_{jt} + \xi_{jt} + \overline{\epsilon}_{ijt}, \tag{3.1}$$

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-time fixed effects as well as retailer age and the variety of products offered by the retailer. 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 econometrician.

To allow for heterogeneity in individual preferences, we model consumer utility over price

[17] We do not consider any quantity choice by consumers as in Dube (2004). Since we do not have consumer-level data, we are not able to estimate such a model.

as

$$\alpha_i^* = \alpha + \Sigma \nu_{ij},\, \nu_{ij} \sim N(0, I_{n+1}), \tag{3.2}$$

where Σ captures the covariance in unobserved preferences over prices. 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. We allow 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 $\overline{\epsilon}_{ijt}$ follows the nested logit distribution, where products in the same group have correlated preferences. We can therefore write this term as:

$$\overline{\epsilon}_{ijt} = \zeta_{jgt} + (1-\rho)\epsilon_{ijt}, \tag{3.3}$$

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, consumers view each product in each nest as perfect substitutes, which in this case implies they have no preference over product type, only at which retailer to shop. Plugging this expression into equation 3.1 gives

$$u_{ijt} = x_{jt}\beta + \alpha_i^* p_{jt} + \xi_{jt} + \Sigma_{g\in G}\chi(j \in g)\zeta_{jgt} + (1-\rho)\epsilon_{ijt}, \tag{3.4}$$

where $\chi(j \in g)$ is a dummy variable indicating if product j is in group g, meaning sold at retailer g. Allowing for a random coefficient on price and a flexible nesting parameter on product type allows for robust substitution patterns. When $\Sigma = 0$ 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 as

$$\delta_{jt} = x_{jt}\beta + \alpha p_{jt} + \xi_{jt}, \tag{3.5}$$

this utility produces market shares:

$$s_{jt}(\delta_{jt}, \theta, \nu_i) = M \cdot \frac{\exp((\delta_{jt} + p_{jt}\Sigma\nu_i)/(1=\rho))}{\exp(I_{ig}/(1-\rho))} \frac{\exp(I_{ig}/(1-\rho))}{\exp(I_i)} \tag{3.6}$$

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

$$I_{ig} = (1-\rho)\log\Sigma_{j\in G}\exp((\delta_{jt} + p_{jt}\Sigma\nu_i)/(1-\rho)) \text{ and} \tag{3.7}$$

$$I_i = \log(1 + \Sigma_g \exp(I_{ig})) \tag{3.8}$$

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.[18] 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. 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.[19] Since the product is perishable, we ignore the potential for consumer stockpiling across months.[20]

[18]In 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.

[19]Different notions of market size have been tested and none of the results that follow are sensitive to this assumption.

[20]Another 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.

3.2 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 $\theta = (\alpha, \Sigma, \rho)$. Formally the GMM estimator is formed from the population moment condition $E[Z' \cdot \omega(\theta)] = 0$. The GMM estimate is

$$\hat{\theta} = \min_{\theta} \omega(\theta)' Z A^{-1} Z' \omega(\theta) \tag{3.9}$$

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 fixed effects with instruments Z then produces a residual term that is equivalent to $\omega(\theta)$.

After including product type, time, retailer and retailer-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 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.[21] 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 addition to potential endogeneity of prices, Berry and Haile (2014) and others note that

[21] These are similar in nature to "Hausman" instruments, which 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.

the heterogeneity terms introduce additional endogeneity into the estimation. In our RCNL specification, this means additional instruments are needed to ensure identification of Σ and ρ, the standard deviation of price preferences and the nesting parameter. To identify ρ 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 three types of instruments, the number of product types sold by the retailer in each month, the average prices of competing products within the retailer, 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 price and 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.

3.3 Results of Demand Estimation

Results from this estimation are shown in Table 5. Two versions of our preferred specification are shown. Column (1) shows results when retailer-month specific fixed effects are included and column (2) shows results without these fixed effects but with retailer characteristics, namely age in months and log of variety of products offered in a given month. This is calculated by summing the number of unique inventory items sold in each month. In both 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 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 consumers decide which retailer to purchase from and then compare products at that retailer rather than choosing a product first and then comparing retailers.

Table 6 shows how estimates of the price coefficient and average own-price elasticity vary across specifications. The average own-price elasticity in our preferred specification, shown in column (3), is -2.71. This is close to the average own-price elasticity across consumer packaged goods of -2.62 (Hanssens (2015)). As expected, the simple logit demand model produces more elastic estimates. Table 7 shows results for the nested logit model with and without processor intercepts included, where in the second category product is defined at the retailer-processor-category level. The two specifications produce very similar median elasticities. For the RCNL models, for different types of fixed effects and instruments the elasticity varies but only within the range of 2.5-3. When the wholesale price instruments are dropped from the estimation, both in the RCNL and simple Logit specifications, the resulting price elasticities are closer to zero, suggesting they do correct some remaining endogeneity in prices.

For our preferred specification in Column (3), which includes all fixed effects and instruments, we also calculate the total elasticity for the marijuana category as a whole relative to the outside good. We find the category is more inelastic, with an aggregate elasticity of -1.47. This suggests most substitution takes place within the marijuana category with only modest substitution to the outside good. By comparison, Miller and Weinberg (2017) find a category elasticity of .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. With 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. Nevertheless we find that demand is relatively inelastic for the category as a whole, suggesting there is not a widely available black market where consumers may find substitute products. 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.

Table 5: Demand Estimates

		(1) RCNL-1	(2) RCNL-2
Price	α	-.090	-.139
		(.003)	(.003)
Usable Marijuana	β	1.230	1.286
		(.405)	(.142)
Solid Edible	β	..383	.424
		(.016)	(.019)
Liquid Edible	β	.460	.573
		(.016)	(.019)
Extract	β	1.443	1.469
		(.041)	(.028)
Retailer Age (months)	β		.039
			(.013)
log(Retailer Variety)	β		.897
			(.339)
Random Coeff. on Price	Σ	.029	.053
		(.0)	(.001)
Nesting Parameter	ρ	.689	.689
		(.0)	(.001)
Type FE		Yes	Yes
Time FE		Yes	Yes
Retailer FE		Yes	Yes
Retailer*Time FE		Yes	

Note: This table presents estimates from the RCNL demand system for two specifications. Product characteristics are price and dummies for type, date and retailer. IV estimation is done using GMM with instruments constructed from wholesale prices, average producer prices, and lagged temperature and rainfall at the producer level. There are 32,939 observations at the type-retailer-month-year level.

Table 6: Price Elasticity Estimates

		(1) Logit	(2) Logit	(3) RCNL-1	(4) RCNL-2	(5) RCNL-3	(6) RCNL-4
Price	α	-.161	-.109	-.090	-.135	-.123	-.067
		(.003)	(.006)	(.003)	(.003)	(.003)	(.003)
Random Coeff. on Price	Σ			.029	.053	.053	.029
Median Own-Price Elasticity		-3.248	-2.518	-2.707	-2.944	-2.536	-1.862
Aggregate Marijuana Elasticity				-1.465	-1.609	-1.378	-1.002
Type FE		Yes	Yes	Yes	Yes	Yes	Yes
Time FE		Yes	Yes	Yes	Yes	Yes	Yes
Retailer FE		Yes	Yes	Yes	Yes	Yes	Yes
Retailer*Time FE		Yes	Yes	Yes		Yes	
Retailer Characteristics					Yes		Yes
Wholesale Price IVs		Yes		Yes	Yes		
Producer Price IV		Yes	Yes	Yes	Yes	Yes	Yes
Weather IVs		Yes	Yes	Yes	Yes	Yes	Yes
First-Stage F-Stat		472.8	175.3	427.2	472.8	175.3	113.3

Note: This table presents price coefficient and elasticity estimates from various specifications. There are 32,939 observations at the type-retailer-month-year level.

3.4 Pass-Through Rate

A key empirical measure of firm conduct is the pass-through rate. Because we directly observe wholesale prices at the transaction level, measuring pass-through is straightforward. This will also be the key input to the framework developed by Fabinger and Weyl (2013), who characterizes how tax incidence and the social cost of taxation under imperfect competition are affected by pass-through and a conduct parameter. In this section, we discuss estimation of the pass-through rate of wholesale prices under different specifications.

Wholesale prices are typically estimated from the assumed supply-side first-order conditions (see, e.g., Conlon and Rao (2016)) but an advantage of our data is that we can directly observe them. Using these data, we estimate the following model to obtain the own pass-

Table 7: Demand Robustness

		(1) Nested Logit	(2) Nested Logit
Median Own-Price Elasticity		-3.250	-3.199
Price	α	-.090	-.132
		(.003)	(.035)
Usable Marijuana	β	1.230	1.321
		(.405)	(.127)
Solid Edible	β	..383	.062
		(.016)	(.227)
Liquid Edible	β	.460	.821
		(.016)	(.137)
Extract	β	1.443	1.167
		(.041)	(.349)
Nesting Parameter	ρ	.909	.552
		(.272)	(.137)
Processor FE			Yes
Type FE		Yes	Yes
Time FE		Yes	Yes
Retailer FE		Yes	Yes
Retailer*Time FE		Yes	Yes
Wholesale Price IVs		Yes	Yes
Producer Price IV		Yes	Yes
Weather IVs		Yes	Yes
Observations		33,175	286,042

Note: This table presents estimates from the Nested Logit demand system for specifications with and without processor fixed effects. Product is thus defined as either retailer-category-month or retailer-processor-category-month. IV estimation is done using GMM with instruments constructed from wholesale prices averaged across processor-regions, and average producer prices and lagged temperature and rainfall at the producer level.

Table 8: Pass-through Estimates: Weekly Average Prices

	(1) OLS	(2) OLS	(3) OLS	(4) IV	(5) IV	(6) IV
Wholesale Price	1.993***	1.992***	1.971***	1.598***	1.595***	1.530***
	(0.00562)	(0.00568)	(0.00570)	(0.0439)	(0.0442)	(0.0654)
Avg. Competitor Wholesale			0.112***			0.153***
			(0.00878)			(0.0143)
Year x Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Product Type FE	No	Yes	Yes	No	Yes	Yes
Observations	125,967	125,959	95,805	125,959	125,959	95,785

Note: The table shows the pass-through estimates of the regression 3.10 with weekly average prices.

through rate.

$$p_{ijt} = \beta_0 + \beta_1 w_{ijt} + \beta_2 w_{-ijt} + x'_{it}\beta_3 + \mu_i + \mu_j + \mu_t + \varepsilon_{ijt}, \tag{3.10}$$

where p_{ijt} is the tax-inclusive weekly-average retail price by retailer i for category j at week t, w_{ijt} is the average wholesale price that retailer i pays for category j at week t, w_{-ijt} is the average wholesale price that competitors pay for category j at week t, x_{it} is a vector of variables for observed retailer characteristics, μ_i 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. Note that there are four different types of marijuana products: $j \in$ {usable, solid edible, liquid edible, and extract}.[22]

Table 8 shows the results of both panel linear and panel IV regression. We use IV regression to take care of potential endogeneity in wholesale prices. Even though our empirical model controls for a rich set of fixed effects, there still might exist some remaining unobserved het-

[22] In principle, the state's tracking system allows us to match retail prices and wholesale prices at the transaction level. We do not estimate the transaction-level pass-through rates because daily-level price variations are noisy and pass-through estimates are not stable.

erogeneity such as the bargaining power of retailers and wholesalers. The set of instruments we exploit for dealing with the endogeneity concern consists of the weather-related variables used in the previous section, namely temperature and precipitation, and average transaction prices further upstream between producers and processors. The validity of these instruments rests on the assumption that weather conditions at the producer's location is likely to effect wholesale prices through the quality of products, but is not likely to directly change retail prices when they are sold. Similar logic applies to the upstream prices, as producer prices are likely to affect wholesale prices but not retail prices directly. These instruments have been used in the prior literature on estimating pass-through rates (see, e.g., Nakamura and Zerom (2010)).

The results show that own pass-through rates are significantly higher than 1 and are actually nearly 2 if wholesale prices are not instrumented. Hence, we find that cannabis retailers pass through their cost shocks more than perfectly. When we use IV estimation, the pass-through is about 1.5. In columns 3 and 6, we include the average wholesale prices of competitors and find a negative association with own retail price, but the magnitude is smaller than the effect of own wholesale price. As discussed in Miravete, Thurk, and Seim (2018a) and Fabinger and Weyl (2013), pass-through greater than 1 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 (2016) and Conlon and Rao (2016). In those studies, the authors find significant market power of retailers in the cement industry and the liquor industry, respectively.[23]

A concern one might have in our pass-through estimates is auto-correlation of the error terms, ε_{ijt}. In the previous specification, we run a series of fixed effect models to control for unobserved heterogeneity, while we assume ε_{ijt} is i.i.d. To see the robustness of our results to

[23] In the estimation, we use tax-inclusive prices to estimate pass-through. When we use tax-exclusive prices, the estimated pass-through rate becomes lower around 1.2. In Conlon and Rao (2016), they report the pass-through rate with tax-exclusive prices.

Table 9: Pass-through Estimates: First-Difference Estimator

	(1) OLS	(2) OLS	(3) OLS	(4) IV	(5) IV	(6) IV
Wholesale price	1.900***	1.901***	1.882***	1.734***	1.753***	1.651***
	(0.00473)	(0.00484)	(0.00506)	(0.0546)	(0.0537)	(0.0538)
Avg. Competitor Wholesale			0.0180**			0.0598***
			(0.00639)			(0.0115)
Observations	124877	124870	94160	124838	94213	94144
R-squared	0.57	0.57	0.64	0.56	0.61	0.61

Note: The table shows the pass-through estimates of the regression 3.11 with weekly average prices. License, category and time fixed effects are included in the estimation.

this concern, we consider the following specification.

$$\triangle p_{ijt} = \beta_0 + \triangle\beta_1 w_{ijt} + \triangle\beta_2 \bar{w}_{-it} + \triangle x'_{it}\beta_3 + \mu_i + \mu_j + \mu_t + \epsilon_{ijt}, \tag{3.11}$$

where $\triangle p_{ijt} = p_{ijt} - p_{ijt-1}$, $\triangle w_{ijt} = w_{ijt} - w_{ijt-1}$. Other variables $\triangle \bar{w}_{ijt}$ and $\triangle x_{ijt}$ are similarly defined.

Table 9 reports the estimation results. Similar to Table 8, we estimate both panel-linear models and IV models and confirm the results we find in Table 8. We find that the own pass-through estimates are still greater than unity for both specifications and the average wholesale prices of competitors has a positive impact, but the coefficient is much smaller than the own pass-through rate.[24]

In Table 10, we estimate the pass-through rate by product category. We find the pass-through greater than unity for all categories, particularly higher for liquid marijuana products. Hence, the results indicate that retailers have market power regardless of the product category.

[24]Taking first difference of equation 3.10, the license fixed effect and the product-type fixed effect are canceled out. In the main text, the models keep these fixed effects, following the pass-through literature, but we also estimate the models without these fixed effects. The results are qualitatively similar and available from the authors upon request.

Table 10: Pass-through Estimates by Category

	(1) Solid	(2) Liquid	(3) Extract	(4) Usable
Wholesale price	1.318***	2.637***	1.321***	1.207**
	(0.0747)	(0.0477)	(0.269)	(0.404)
Observations	31372	28494	31278	33662
R-squared	0.58	0.63	0.26	0.51

Note: The table shows the pass-through estimates of the regression 3.11 with weekly average prices. License, category and time fixed effects are included.

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 using the weekly aggregated data as in Table 8 and report the results in Table 5. We find that there is a general downward trend over time. In 2015, the average own pass-through rates were above 2, but they eventually get smaller to about 1.9 in 2017. This trend may imply that the retail cannabis market is becoming more competitive over time.

One may wonder why pass-through is so high in the Washington marijuana market. One potential reason would be discrete prices as discussed in Conlon and Rao (2016). They find that 77% of price changes in the distilled spirits market in Connecticut are in whole-dollar increments and it leads to excessive pass-through. We investigate this possibility in our data. We find no evidence of 9-ending prices in retail prices nor wholesale prices. In addition, price changes are not $1, but look more continuous for both positive and negative price changes.

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 some market power and more tax burden falls on consumers than retailers. These results also strengthen the conclusion that there is not readily available black market

Figure 5: Monthly Pass-Through Rate of Wholesale Prices on Retailer Prices

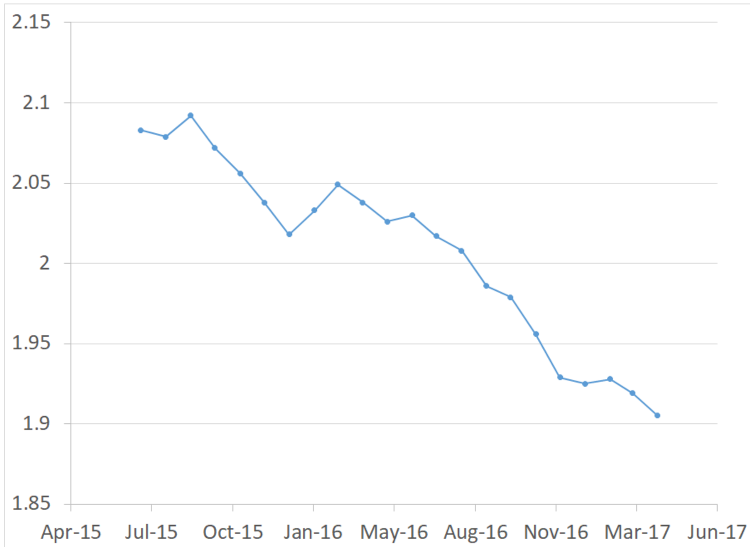

marijuana acting as a substitute for legal marijuana sales. If this black market existed, retailers would not be able to pass-through their cost-shocks more than fully without losing excessive sales. In the next section, we employ the framework developed by Fabinger and Weyl (2013) to quantify the incidence of Washington's excise taxes by combining the pass-through estimates and the consumer demand estimates.

4 Policy Analysis

In this section, we use our empirical results to examine how to regulate the recreational cannabis industry. We begin by calculating 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 show that the state is on the upward sloping region of the Laffer curve. Finally, we impose this supply side model of firm behavior and examine how much additional tax revenue the state could earn with higher taxes, as well as what effect these would have on total consumption. We also use this model to evaluate other regulatory policies including a state monopoly on marijuana sales. In each case, we highlight the effect of retailer market power on these outcomes.

4.1 Policy Analysis: Tax Incidence

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 $I = \frac{\rho}{1-\rho}$ where ρ 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:

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

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. In the case of the Washington marijuana industry, the state sets a cap on the total amount of production and can set this cap to bind in equilibrium. Thus supply is likely to be perfectly inelastic and consumers will pay the entire tax with no deadweight loss associated with taxation.

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

$$\frac{p-mc}{p}\epsilon_D = \theta, \tag{4.1}$$

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:

$$\frac{dPS}{dt} = -[1-\rho(1-\theta)]q \tag{4.2}$$

and the marginal effect on consumers is

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

Thus, the tax incidence can be calculcate as

$$I = \frac{\rho}{1-\rho(1-\theta)}, \tag{4.3}$$

where in this case

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

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 4.1. A more complete description of this counterfactual is described in the following section. We estimate an average $\hat{\theta} = .79$, in other words observed margins are 79% of the hypothetical monopolist's margins, with a 95% confidence interval of $(.61, .99)$.

Equation 4.3 gives the ratio of consumer harm to producer harm from a small unit tax increase.[25] Using estimated $\rho = 1.9$ and $\theta = .79$, implied incidence of taxes falls roughly 27% on producers and 73% 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 $5,000$kg 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 \$9.0 million while producer profits would fall by \$3.1 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.

[25]While in practice Washington uses ad valorum 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 valorum tax on retail sales. In addition, several other states including California do impose unit taxes.

[26]If we use $\rho = 1.6$ and $\theta = .79$, then consumers bear about 71% of the tax burden. These results are along the lines of what Conlon and Rao (2016) find in the liquor industry, in which consumers bear between 75% to 80% of the tax burden.

4.2 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, Thurk, 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.

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

$$\frac{\partial \pi_r}{\partial p^r} = (p^r - p^w)\frac{\partial D((1+\tau)p^r)}{\partial p^r}(1+\tau) + D((1+\tau)p^r) \tag{4.4}$$

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

Applying the Implicit Function Theorem to equation 4.4, the tax pass-through rate can be written as

$$\frac{dp^r}{d\tau} = \frac{\kappa(p^*) - (2 - \frac{p^w}{p^r})}{p^*(2 - \kappa(p^*))}, \tag{4.5}$$

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

$$\eta(\tau) = \frac{\partial p^r}{\partial \tau} \times \frac{\tau}{p^r} = -\frac{\tau}{1+\tau} \times \frac{(1-\frac{1}{\varepsilon(p^*)}) - \kappa(p^*)}{2-\kappa(p^*)}. \tag{4.6}$$

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

$$\frac{dp^r}{dp^w} = \frac{1}{2-\frac{DD''}{D'^2}} = \frac{1}{2-\kappa(p^*)}, \tag{4.7}$$

which can be written as a function of the demand curvature, κ. Combining equations (4.6) and (4.7), we can rewrite the tax elasticity as

$$\eta(\tau) = -\frac{\tau}{1+\tau} \times [(1-\frac{1}{\varepsilon(p^*)}) - (2-\frac{1}{\frac{dp^r}{dp^w}})] \times \frac{dp^r}{dp^w}, \tag{4.8}$$

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

$$R'(\tau) = p^r D(p^*)\left[1 + \frac{\tau}{1+\tau}\varepsilon(p^*) + \eta(\tau)(1+\varepsilon(p^*)\right] = 0. \tag{4.9}$$

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 4.6 and 4.9 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 4.7 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 (2001), $\eta(\tau) > 0$ 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 -1.14 as shown in Table 6, the industry would clearly be on the upward sloping portion of the Laffer curve in the monopoly case.

Multi-product Oligopoly We now consider the more general case with multiple asymmetric retailers, each selling multiple products. With J retailers and K manufacturers transacting L products. Retailer i's profit function is

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

where j denotes product, J_i denotes the set of products that retailer i sells, p_j^r is retailer price of product j charged by retailer i, p_j^w is wholesale price of product j paid by retailer i, and p is a $J \times 1$ vector of retail prices $\{p_j^r\}$. In Appendix A 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 4.9. 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.

4.3 Counterfactual Policy Simulations

In this section we study how much more revenue the state could earn with higher taxes and what would the effect be on total consumption and retailer profits. We also consider alternative regulatory arrangements including a state monopoly on retail sales. 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. This is a standard assumption in industrial organization, and typically uses estimated marginal costs in addition to estimated demand.

To evaluate the fit of this model we compare its predicted pass-through to observed pass-through. We solve for the equilibrium prices under observed wholesale prices and then simulate a small cost shock to measure the amount of equilibrium pass-through. Under Nash-Bertrand oligopoly competition and our estimated demand model, we get an average pass-through rate of 1.75, very close to observed pass-through rates.

Table 11: Counterfactual Tax Policy (2017)

	15%	37%	40%	50%
Monthly Tax Revenue (millions of $)				
Strategic Price Reaction	12.3	25.3	26.7	30.8
Fixed Prices	11.70	25.3	26.5	29.2
Average Pre-Tax Price				
Strategic Price Reaction	20.82	18.63	18.51	18.13
Fixed Prices	18.63	18.63	18.63	18.63
Consumer Welfare				
Strategic Price Reaction	46.4	40.2	39.3	36.3
Fixed Prices	57.3	40.2	38.3	32.8

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.

Tax Policy Counterfactual 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 11 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 $1.4 million per month, a 5.5% increase. Increasing the tax rate to 50% would increase tax revenue by $5.5 million per month, a 21.7% increase. We also find that increasing the tax rate to 40% would cause retail sales of usable marijuana to fall by approximately 25kg, or 1.2% of total sales.[27] Because prices would be out of sample under further increases in the hypothetical tax, we choose not to extrapolate out for tax rates higher than 50%.

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 $1.6 million fewer dollars per month, or a 29% 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 off by nearly one third.

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 $12.3 million, less than half of its current revenue. On an annual basis this would amount to $156 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.[28]

[27] This 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.

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

Table 12: Market Structure Counterfactuals (2017)

	Current Policy	Single-category Retailers	State Monopoly No Tax	State Monopoly	51.9% Margin
Avg. Pre-tax Price ($)	18.63	11.62	22.79	19.36	11.83
Tax Revenue (millions of $)	25.3	29.79	0	24.94	27.60
Usable Sales (kg)	2109	3301	2297	2023	4376
Retailer Revenue (millions of $)	68.5	80.5	99.5	67.4	74.6
Retailer Variable Profits (millions of $)	40.4	29.9	64.3	40.6	25.5
Consumer Surplus (millions of $)	20.1	43.3	26.2	19.5	45.8

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.

Market Structure Counterfactual 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 statewide 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 Table 12. 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 as well. We find that tax revenue would increase by 15%.

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 would charge the profit maximizing price. Results from this counterfactual are shown in column 3 and 4 of Table 12. 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 and the state simply earns the retail profits as revenue. With the 37% tax still in place, we find the state retailer would increase prices, but only slightly. Total sales would fall slightly as would tax revenue. The state would earn $66 million per month from combined tax revenue and retail profits. If the state eliminated the tax as a source of revenue but kept all retail profits under the state monopoly, this would by slightly lower, at $64 million in monthly 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%. This follows from the fact that consumers do not appear to search actively and the private retailers already behave as local monopolists.

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 12 and find that prices would fall significantly, total sales would increase and total tax revenue would also substantially increase.

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.

5 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 data to investigate the incidence of these taxes and whether the state is overtaxing the product and reducing revenue, as well as the role of market power in designing the retail sales tax scheme. 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 (1994). 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, roughly 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*.

ANDERSON, S., A. DE PALMA, AND B. KREIDER (2001): "Tax Incidence in Differentiated Product Oligopoly," *Journal of Public Economics*.

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 BerkeleyâĂŹs 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).

CONLON, C., AND N. RAO (2016): "Discrete Prices and the Incidence and Efficiency of Excise Taxes," *Working Paper*.

DRAGONE, D., G. PRAROLO, P. VANIN, AND G. ZANELLA (2017): "Crime and the Legalization of Recreational Marijuana," *mimeo*.

DUBE, J.-P. (2004): "Multiple Discreteness and Product Differentiation: Demand for Carbonated Soft Drinks," *Marketing Science*, 23(1), 66–81.

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*.

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.

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, N., M. OSBORNE, AND G. SHEU (2016): "Pass-Through in a Concentrated Industry: Empirical Evidence and Regulatory Implications," *RAND Journal of Economics*.

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*.

NAKAMURA, E., AND D. ZEROM (2010): "Accounting for Incomplete Pass-Through," *The Review of Economic Studies*, 77(3), 1192–1230.

NEVO, A. (2001): "Measuring Market Power in the Ready-to-Eat Cereal Industry," *Econometrica*, 69, 307–342.

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.

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 The Laffer Curve under Multi-product Oligopoly

In this appendix, we extend the results presented in section 4.2 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.

With J retailers and K manufacturers transacting L products. Retailer i's profit function is

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

where j denotes product, J_i denotes the set of products that retailer i sells, p_j^r is retailer price of product j charged by retailer i, p_j^w is wholesale price of product j paid by retailer i, 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

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

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 (A.1) gives the (own) pass-through rate of wholesale prices to retail prices as follows:

$$\frac{dp_j^r}{dp_j^w} = \frac{\frac{\partial D_j}{\partial p_j^r}}{2\frac{\partial D_j}{\partial p_j^r} + \sum_{j' \in J_i} \left[(p_{j'}^r - p_{j'}^w)(1+\tau) \frac{\partial^2 D_{j'}}{\partial p_j^{r2}} \right]}. \tag{A.2}$$

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

$$\frac{dp_j^r}{d\tau} = \frac{\sum_k \frac{\partial D_j}{\partial p_k^r} p_k^r + \sum_{j'} \left[(p_{j'}^r - p_{j'}^w) \left((1+\tau) \sum_k \frac{\partial \left(\frac{\partial D_{j'}}{\partial p_j^r} \right)}{\partial p_k^r} p_k^r + \frac{\partial D_{j'}}{\partial p_j^r} \right) \right]}{2\frac{\partial D_j}{\partial p_j^r} + \sum_{j' \in J_i} \left[(p_{j'}^r - p_{j'}^w)(1+\tau) \frac{\partial^2 D_{j'}}{\partial p_j^{r2}} \right]} \tag{A.3}$$

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

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

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

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

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 A.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 4.9. 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.