The Taxation of Recreational Marijuana: Evidence from Washington State

Benjamin Hansen*†
University of Oregon, NBER, IZA

Keaton Miller*
University of Oregon

Caroline Weber*
University of Oregon

August 2017

Abstract

More than half of American voters support marijuana legalization, in part to increase tax revenues. States with legal marijuana markets have wildly different taxes and regulations, indicating diverse policymaker beliefs. We examine a Washington policy reform: a switch from a 25% gross receipts tax collected at every step in the supply chain to a sole 37% excise tax at retail. Using comprehensive administrative data, we assess market responses. We find that the previous tax regime strongly incentivized vertical integration. Consumers bear 44% of the tax burden. Consumer marijuana demand is price-inelastic in the short run, but becomes price-elastic within a few weeks.

JEL Codes: H20, H30, H70, I10, K40.

Keywords: Marijuana, Excise Taxes, Pass Through, Tax Incidence, Vertical Integration, Tax Invariance, Natural Experiment

*Address: 1285 University of Oregon, 435 PLC, Eugene OR, 97403-1285. Phone: 541-346-8845. Hansen: bchansen@uoregon.edu; Miller: keatonm@uoregon.edu; Weber: cweber5@uoregon.edu

†The authors would like to thank David Agrawal, Nathan Anderson, David Evans, Michael Grossman, Donald Kenkel, Michael Kuhn, Nathan Seegert, and Joel Slemrod for helpful comments. We appreciate comments and feedback from participants at seminars at Cornell, Columbia, Rutgers, Case Western, Portland State, and conference participants at IHEA, NTA, WEAI, and the IIOC meetings, as well as industry participants and Cannabis Science and Policy Summit attendees. Kendall Houghton provided exemplary research assistance.

1 Introduction

The United States reached a tipping point on marijuana in 2015, when, for the first time, the majority of adults indicated support for legalizing marijuana in some form (Motel 2015). Four states, Alaska, Colorado, Oregon, and Washington, currently have legally operating recreational marijuana markets. During the 2016 elections, residents of California, Maine, Massachusetts, and Nevada voted for legalization.[1] Internationally, Canada, Colombia, the Netherlands, Spain, South Africa, and Uruguay have also legalized marijuana for recreational use in some form.

Taxation and revenue generation has been one of the primary arguments for legalizing recreational marijuana (Miron and Zwiebel 1995). However, little is known about the behavior of participants in legal marijuana markets, and therefore what is its revenue potential. In particular, there is little evidence on how industry participants may respond to alternative policies - particularly given the available substitutes in the form of black-market goods and legal or quasi-legal medical marijuana. Policymakers across states have implemented varied regulatory regimes: Colorado requires retail firms to grow 80% of the product they sell, while Washington forbids retailers from growing at all. Washington's current tax rate is 37%; Maine has a tax rate of 10%.[2] States also apply taxes at different points in the supply chain. These differences point to substantial variation in beliefs among policymakers and a need for more information about the impact of different policy levers.

We offer new evidence concerning key policy-relevant behaviors by examining the market in Washington where, after a 2012 vote, recreational marijuana became available at retail stores beginning on July 8, 2014. Washington's regulations created a supply chain with three types of firms: cultivators, who grow the cannabis plant; processors, who transform the raw plant material into usable marijuana[3], and retailers, who sell products to end-users.

[1] Though marijuana is considered a Schedule I Controlled Substance by the federal government, the Department of Justice has not enforced prohibition in the relevant states.

[2] See Table 1 for detail on state-level tax rates.

[3] Here and throughout, we discuss only 'usable marijuana', which consists of the dried and cured flowers of the cannabis plant, and is consumed by smoking. Different 'products' consist of different strains of marijuana,

To comply with state law, firms must provide detailed information about their operations to the state government. These unique[4] "seed-to-sale" administrative records track the entire legal supply chain, and are verified through frequent government audits. We have obtained this data, which gives us the unique ability to observe the prices, quality, and variety of marijuana products in the marketplace. In contrast, previous studies of the marijuana industry have examined the illegal market through surveys, prices collected from drug seizures, and crowd-sourced data (Miron and Zwiebel 1995; Pacula et al. 2000; Clements and Zhao 2009; Donohue, Ewing, and Peloquin 2011; Williams, van Ours, and Grossman 2011; Jacobi and Sovinsky 2016). We are the first to examine a legal market directly.

We take advantage of a reform of Washington's tax policy. Prior to July 1, 2015, a 25% gross receipts tax was assessed at each transfer of marijuana (i.e., between cultivators and processors, processors and retailers, and retailers and consumers). After July 1, 2015, the only tax collected was a 37% excise tax at retail. Crucially, this change was unexpected by market participants: the reform was passed during a special session of the Washington Legislature on June 27, 2015, and signed by the Governor on June 30. Contemporaneous media reports suggest that, although the change was supported by the industry, it was not expected to pass, and market participants did not have confidence in their forecasts of future prices at the time of the change (La Corte 2015). We estimate the effect of this exogenous change on key decisions made throughout the production process, while controlling for product quality and other factors that may impact demand and supply.

We find that the original tax regime encouraged vertical integration between cultivators and processors (Washington bans vertical integration between retailers and other firm types). We define vertically integrated marijuana as product for which the cultivator and processor were the same firm. We capture changes in vertical integration by examining two measures:

which are analogous to the different cultivars of apples commonly available in supermarkets. While other derivatives of the cannabis plant are available, usable marijuana comprises the vast majority of the market by transaction count and revenue.

[4]While Colorado, which legalized marijuana at roughly the same time as Washington, maintains a similar set of administrative records, they are legally proprietary.

first, we consider the fraction of wholesale transactions to retailers of vertically integrated marijuana, and second, we consider the total weight of vertically dis-integrated marijuana sold to retailers. Before the reform, roughly 94% of wholesale transactions comprised vertically integrated marijuana. This drops to about 90% following the tax change. The weight of dis-integrated marijuana sold to retailers increased by 42% after the change.

Taken together, these results suggest that the original gross receipts tax created inefficiencies in the supply chain that persisted after the removal of the taxes. To test this hypothesis, we restrict the analysis to transactions stemming from producers in their first week of business—i.e., to capture firms' decisions about vertical integration before paying the fixed costs associated with entry. We find that the drop in the percentage of vertically integrated transactions quadruples to 16%, and the increase in the weight of dis-integrated marijuana more than doubles to 105%. This provides the first empirical evidence in the literature that gross receipts taxes lead to vertical integration.

The tax-invariance folk theorem makes the claim that if taxes remain the same, but a tax reform changes where in the supply chain they are assessed, after-tax prices at each point in the supply chain will remain the same. Because we find that the tax reform eliminated 28 cents of taxes overall, but the remaining 75 cents of the processor tax were shifted from the processor to the retailer, the tax-invariance folk theorem suggests that processors would lower their wholesale prices by 75 cents. However, we find that processors only lower their prices by an average of 26 cents per gram, suggesting that the tax invariance folk theorem has been violated in this setting.

While the price elasticity of demand and tax incidence are of first-order concern for policymakers tasked with setting taxes, the analysis is complicated by the degree of product differentiation—the typical retailer has many different products available at any time. Given the variation in product characteristics, we test for and find a significant decrease in the potency of marijuana purchased by consumers, as measured by the concentration of the psychoactive chemical tetrahydrocannabinol (THC), after the reform. We find no change in

the potency of marijuana grown or sold at wholesale near the tax change. Taken together, these results suggest that consumers reacted to the reform in part by substituting toward lower quality products.

We account for this substitution by estimating the retail price response with a specification that includes fixed effects for the interaction between cultivator, processor, retailer, and product. We find that the tax-inclusive price faced by consumers for identical products increased by 2.3%. Considering the changes in tax rates as well as the adjustment in retailers' marginal costs due to the change in processor prices in response to their own tax changes, we find that consumers bear about 44% of the retail tax burden. We find that the quantity purchased decreased by 0.95%, though estimates are noisy, implying a short-term price elasticity of -0.43. However, over time, the magnitude of the quantity response significantly increases, and our estimates suggest that the price elasticity of demand is about negative one within two weeks of the reform. We conclude that Washington, the state with the highest marijuana taxes in the country, is near the peak of the Laffer curve: further increases tax rates may not increase revenue.

Our vertical integration findings contribute to long-standing broad public finance questions. It is theoretically clear that vertical integration would be a natural consequence of a gross receipts tax, much more so than other taxes which may encourage vertical integration, such as a value added tax (Kopczuk and Slemrod 2006). Tax economists frequently come out vehemently against gross receipts taxes for this and other reasons (e.g., McLure 2005; Pogue 2007; Testa and Mattoon 2007); however, this is the first article we are aware of that provides compelling empirical evidence of this behavior. Gross receipts taxes have begun to proliferate across states in recent years (Kaeding 2017), so this article provides an important source of empirical evidence that such taxes do, in fact, lead to inefficient levels of vertical integration.

We also contribute to the discussion on tax invariance along two dimensions. First, we highlight that the possibility of vertical integration under a gross receipts tax is another

reason tax invariance will no longer be expected to hold. Second, we document that tax collection at the retail level has different incidence implications than it does at the processor level. In contrast to previous work, there is no evidence that the result is driven by different tax avoidance or evasion possibilities at different points in the supply or consumption chain (Slemrod 2008; Kopczuk et al. 2016; Brockmeyer and Hernandez 2016), nor is it likely driven by a relative lack of awareness of the tax by one side of the market (Chetty, Looney, and Kroft 2009). Instead, media reports and conversations with industry participants suggest that processors took advantage of a unique opportunity to increase margins.

Our incidence results stand in contrast to the literatures on cigarette and gasoline tax incidence, which generally find that consumers bear a substantial majority of the tax burden (Harding, Liebtag, and Lovenheim 2012; Kopczuk et al. 2016) and these pass-through rates can even exceed one (Barnett, Keeler, and Hu 1996; Kenkel 2005). One plausible explanation for the difference is that our estimated medium-run elasticity of demand for marijuana is higher than the consensus estimates for cigarettes or gasoline. Alternatively, tight ownership and size restrictions in Washington's marijuana market may lead to differential market power effects or frictions relative to these other markets. This difference could also be explained by the decline in the amount of federal income taxes some retail firms expected to pay upon the switch from a gross receipts tax to an excise tax.[5]

Finally, our findings also contribute to broader research on supply-side interventions in drug markets, which has included examinations of policies ranging from the prohibition of alcohol (Miron and Zwiebel 1991), to limitations on precursors of methamphetamines (Dobkin and Nicosia 2009) and the legalization of the cultivation and sale of marijuana for medical purposes (Anderson, Hansen and Rees 2013). We are the first to study legalized marijuana markets for recreational use, and the first to analyze the substantial natural

[5]The gross receipts tax was defined as part of the price used to calculate other state and local taxes for recreational marijuana retail firms. Hence, some portion of participants in the retail market expected to be substantially better off as a result of the tax regime change because they would no longer owe federal income taxes on their state gross receipts tax. However, on July 31, 2015, the IRS issued a memo clarifying this issue and noted that these firms never did owe federal income taxes on their state gross-receipts tax liability (Moffitt 2015). We provide more details on this in Section 4.3.

experiment that a tax change offers. And while prior studies have focused on the elasticity of demand within illegal marijuana markets, we study the market-level elasticity of demand in a legal marijuana market.

We proceed as follows. In Section 2, we discuss the history of marijuana and the tax system in Washington. In Section 3 we discuss the administrative data we utilize and the methods we use to estimate responses to the policy change. We present our results in Section 4. We conclude in Section 5 with some policy and economic implications of our findings.

2 Background

Prior to 1938, marijuana was legal in the United States. Indeed, it was listed in the United States Pharmacopeia as a prescription for labor pains, nausea, and other conditions. Since the passage of the Marijuana Taxation Act of 1938, the consumption of marijuana has been illegal. The advent of Scheduled substances in the Controlled Substances Act of 1970 significantly strengthened the prohibition against marijuana, as it was quickly classified a Schedule I substance with a 'high potential for abuse and little known medical benefit.'[6]

In 1996, California became the first state to legalize marijuana for medical use. In 1998, Washington voters also legalized medical marijuana through Initiative 692. Currently, 27 states and regions permit the cultivation and use of marijuana for medical reasons. In response to the growing acceptance of medical marijuana, in October 2009, the United States Department of Justice issued a memorandum to United States Attorneys (Ogden 2009) discussing the allocation of resources in states with legal medical marijuana markets. The memo specified that "federal resources in your States" should not be focused "on individuals whose actions are in clear and unambiguous compliance with existing state laws providing for the medical use of marijuana." though the memo also emphasized a need to investigate and prosecute "drug traffickers who hide behind claims of compliance with state law" (Ogden 2009, p. 2). This has been broadly interpreted as an effort to defer to states' choices in the

[6]Other Schedule I substances include heroin and methamphetamine.

absence of federal consensus (Stout and Moore 2009).

In the election of November 2012, Washington voters approved Initiative 502, which legalized the possession and recreational consumption of marijuana for adults over 21.[7] The law created three types of licenses: one for producers (whom we term 'cultivators' in this article), who are permitted to grow and harvest cannabis plants; another for processors, who transform the harvested plant material into usable marijuana and other products for wholesale distribution; and a third for retailers, who may sell the final products that they obtain at wholesale prices to consumers. The Initiative specified a tax structure with a 25% tax on the revenue of each marijuana transaction. This included the sale of grown marijuana to processors who convert the harvested plant material into usable marijuana[8] and other products, the sale by processors of usable marijuana to retailers, and the sale of final products by retailers to end consumers.

Cultivator licenses come with capacity constraints, meaning that each cultivator is allocated one of three sizes of plant canopy[9] Firms that grow cannabis plants may also have a license to process the material into usable marijuana, and vice versa. Firms involved in the production or processing of marijuana, however, are forbidden from owning or operating a retail location. Under the original tax regime, vertically integrated producers and processors did not owe any taxes on internal transfers of marijuana.

In August 2013, the federal Department of Justice responded to the changing environment in Washington and Colorado[10] with a memo commonly known within the industry as

[7]Recreational use of marijuana remained illegal in Oregon, Idaho, and British Columbia, Washington's neighbors, until October 1, 2015, when retail stores in Oregon opened. The effects of legalization in Oregon on Washington's market is discussed by Hansen, Miller, and Weber (2017).

[8]"Usable marijuana" is defined by Washington state law as "dried marijuana flowers, [excluding] marijuana-infused products [and] marijuana concentrates." In practice, usable marijuana is consumed either through the use of a fixed apparatus or by rolling the flower into a "joint" with paper produced for the purpose. As other product categories include different, difficult-to-distinguish levels of value-add at the processor level, we focus on the "usable marijuana" category exclusively.

[9]According to Washington law (WAC 314-55-010), " 'Plant canopy' means the square footage dedicated to live plant production, such as maintaining mother plants, propagating plants from seed to plant tissue, clones, vegetative or flowering area. Plant canopy does not include areas such as space used for the storage of fertilizers, pesticides, or other products, quarantine, office space, etc."

[10]Colorado voters also approved a legalization effort in November, 2012.

the Cole Memo (Cole 2013). As with its response to changing views on medical use, the Department emphasized the prohibition on production and consumption of marijuana under federal law, but provided guidance as to specific enforcement priorities. These included concerns about the diversion of products from the legal market to illegal markets or jurisdictions without legal markets, as well as public health concerns associated with marijuana consumption. Importantly, the Department established a clear expectation that "states and local governments... will implement strong and effective regulatory and enforcement mechanisms." To comply with the enforcement priorities laid out in the Cole Memo, Washington implemented a "traceability" system to track the cultivation, testing, processing, and retail sales of marijuana. We provide more details on the administrative data gathered from this system in Section 3.

The tax reform analyzed in this article was part of House Bill 2136 introduced during the 2015 Regular Session of the Washington Legislature. The reform eliminated the 25% taxes within the supply change and increased the retail tax rate from 25% to 37%. The other regulatory details described above were left unchanged. Because our identification rests on the assumption that the policy change created unanticipated cost shocks throughout the supply chain, the details of the bill's history are critical. Table 2 provides a summarized time line of the bill's progress through the Washington Legislature. The bill originated in the House midway through the 2015 Regular Session and accumulated a number of amendments and substitutions before being passed by the full House. While a Senate committee recommended passage on the last day of the session, however, the full Senate declined to consider the bill during the session. A similar pattern occurred when the bill was reintroduced in the First Special Session: the House quickly passed the bill, and the Senate chose not to take action. Finally, at the very end of the Second Special Session, June 27, the bill received a full Senate vote. The Governor signed it on June 30, and the law went into effect the next day. Contemporaneous media reporting portrayed the industry as unprepared for the change, with one retail store manager quoted as follows: "This is supposed to happen

tomorrow. You have a few hours to change an entire market's pricing structure. It is an exceptionally short window for such a tremendous change" (La Corte 2015).

While the tax change is the most relevant part of the bill for our purposes, the bill, along with companion legislation, contained a number of other measures that arguably may have played a larger role in the internal political process leading to its eventual passage. Increased funds from recreational market tax collections were made available to local jurisdictions, on the condition that they allow firms to enter the market within their geography. Local jurisdictions also obtained greater zoning flexibility for marijuana businesses. Finally, Washington's medical marijuana market, previously legal though essentially unregulated, was brought into the regulatory framework created by the original initiative legalizing recreational use.

Today, eight states have legalized marijuana for recreational use. Table 1 delineates the tax structure within each state. Washington notably applies the highest tax rate of 37%—the next highest is neighboring Oregon, with a 17% tax. Not all states apply taxes at retail—Nevada applies its tax of 15% at wholesale. California applies a cultivation tax designed as a fee for each ounce of dried plant flowers and leaves, while Colorado applies a tax at wholesale based on the average price per gram state-wide. All of these taxes lead to a lower effective rate than Washington's—implying that if Washington is on the left side of the Laffer curve, these other states likely are as well.

In addition to the difference in tax rates and calculation methods, states also differ in regulatory requirements for testing, data collection, zoning, and vertical integration. Washington and Colorado provide extreme examples: Washington bans vertical integration between retailers and all other firms, whereas Colorado requires retailers to grow a large percentage of all product they sell. The administrative data collected by Washington is publicly accessible, whereas Colorado's data is proprietary.

3 Data and Methods

Our data consist of administrative records from the "traceability" (or seed-to-sale) system maintained by the Washington State Liquor and Cannabis Board (WSLCB). The system's purpose is to track each step in the legal marijuana supply chain, enabling regulators to collect taxes and prevent diversion to the black market. Informational reporting compliance is enforced through random audits, backed by penalties that include inventory seizure and destruction. The end result is data that tracks the planting, harvest, and production of cannabis plants into usable goods, the sale of those goods to a retailer, and final retail sale of marijuana products. Along the way, products are tested for potency via the concentration of tetrahydrocannabinol (THC), tetrahydrocannabinolic acid (THC-A), and cannabidiol (CBD), the primary psychoactive components in marijuana.

We utilize an extract of the state's database that includes all plants and products, but removes information that is subject to security or privacy concerns.[11] Firms, locations, and production rooms are given unique identifiers. Each plant is registered at the time of planting. Firms record the provenance of the plant material (e.g., a clone or a seed) as well as the strain.[12] Once harvested, flowers and other plant material are collected and converted into a new "inventory lot" that is assigned a unique identifier (ID); products or material within a single inventory lot are assumed to be homogeneous. These intermediate products may progress through several processing steps before wholesale distribution.

The last processing step is the division of a large wholesale inventory lot into multiple smaller lots for sale to individual retail stores. Each retail lot consists of multiple sealed packages of a specific weight of marijuana (e.g. 1 gram, 2 grams, or 3.5 grams) which are considered identical. When lots are sold to retailers, the tracking system records the date, quantity, and price of the transaction and assigns a new lot ID. Consequently, the observation of a retail lot ID uniquely identifies the retailer, processor, and cultivator, as well as the strain

[11] Our extract does not include data on individual employees or detailed information on transfer logistics such as delivery routes and times.

[12] Strains are defined by the cultivator.

and package size. Transfer manifest identifiers allow us to observe wholesale transactions that comprise multiple product lots. Similarly, the system tracks individual retail transactions, linking the prices and quantity of different items, as well as the transaction time,[13] in a retail transaction to the relevant inventory lots.

We examine both tax-exclusive and tax-inclusive prices. The tax-inclusive prices include both the marijuana-specific excise tax and Washington's general and locality-specific sales taxes. Given the retail firms' inability to use traditional financial services, almost all choose to set tax-exclusive prices that lead to round numbers when taxes are included (this lowers the cash handling costs for firms as they no longer have to handle change). As a consequence, the posted prices faced by consumers include all taxes and, in contrast to many other settings, the sales taxes are therefore salient to consumer-level decision making. The tax reform also changed the way firms reported prices in the traceability system. Prior to the reform, most firms reported fully tax-inclusive prices. Afterwards, most firms reported fully tax-inclusive prices. We clean the price data for each firm to reflect the prices faced by consumers.

We analyze the effects of the tax change on a number of observable behaviors of market participants through a series of regression exercises. We restrict our analysis at the processor and retail levels to the "usable marijuana" product category—74.5% of the total transactions observed in our data.[14] For each component of the supply chain—cultivators, processors, and retailers—we collapse the data by firm-day, unless otherwise specified. We perform minor cleaning steps and exclude firms for which we do not have at least one month of data on either side of the reform. None of this cleaning significantly changes our results. We provide additional details in the Online Data Cleaning Appendix.[15]

We conduct regression analysis to identify responses to the tax change. Because the

[13]Washington does not require retailers to have a constant connection to the tracking system. Many connect to the system at the end of their business day and upload their transactions within one session. To ensure consistency across firms, we choose the daily level as our most granular view of industry activity.

[14]Due to limitations of the traceability system, the "usable marijuana" category we consider contains two types of products: both raw dried flowers and pre-rolled joints, which include some small additional value.

[15]Data cleaning appendix available at `http://www.keatonmiller.org/s/data-cleaning-appendix.pdf`

change took place within the broader context of the market's non-linear evolution, our estimating equations include a polynomial in time. Furthermore, because market participants did not expect the change, we include dummy variables for the days immediately surrounding the tax change, to account for short-term adjustment effects and the Fourth of July, which took place just after the reform was implemented. We analyze behaviors at the firm-day level, though cyclical patterns at the weekly and monthly levels lead us to include additional fixed effects. Our window of analysis spans the two months before and two months after the tax change. We examine the robustness of our estimates to this time window.

Our analyses, therefore, use the following template:

$$\begin{aligned} log(y_{it}) =& \alpha_0 + \alpha_1 tax_{it} + \sum_{j=1}^{6} \alpha_{3j} tax\_day_j + \sum_{k=1}^{6} \alpha_{3k} dow_k + \sum_{l=1}^{31} \alpha_{4l} dom_l \\ &+ \sum_{m=1}^{5} \alpha_{5m} date_{it}^m + \alpha_{6i} + u_{it}. \end{aligned} \tag{1}$$

where y_{it} is our outcome variable for firm i at date t, tax_{it} is an indicator variable for the tax change that is one after July 1, 2015 and zero before, tax_day_j are indicator variables for June 30 - July 5 to absorb local responses to the tax change and the 4th of July, dow_j are day of the week fixed effects, dom_k are day of the month fixed effects, $date_{it}^l$ is a polynomial in the date of sale, and α_{6i} are firm fixed effects. June 30 is the date the law was signed, and July 4th is a holiday that resulted in deviations from normal trends. We explore the sensitivity of our estimates by shrinking our adjustment window to the days June 30 and July 1, or by extending it to June 27th (the date the legislation was passed). We also explore sensitivity to the inclusion of day-of-month fixed effects, as these may not always be very important, but do consume many degrees of freedom. Lastly, we consider sensitivity to changes in the polynomial as well as the inclusion of county-specific polynomials to allow for county-specific time trends.

We take the logs of all outcomes, unless otherwise specified, because these outcomes are essentially log-normally distributed (with the added benefit of allowing us to interpret

the estimated coefficients on the binary regressors as semi-elasticities). Standard errors are clustered by firm. Given our estimating equation, we interpret coefficients on the tax change indicator variable as the average short-term response to the tax reform unless otherwise specified.

Our estimating equations implement a regression discontinuity design, where time is the running variable which determines if firms are exposed to treatment.[16] Hausman and Rapson (2017) recently formalized some of the differences between time-based regression discontinuity applications, which are similar to event studies, and other uses of the regression discontinuity framework. For example, because our data are panel data, testing for covariate balance (Hahn, Todd, and Van der Klaauw, 2001) or smoothness in the running variable (McCrary, 2008), is unnecessary as these features are guaranteed by the panel nature of our data.

4 Results

We report findings in the order marijuana flows through the supply chain: first cultivators, then processors, and finally retailers. Throughout, we present tables with point estimates for each of the outcomes we examine and tables of robustness checks for particularly significant outcomes. For these outcomes, we also provide figures that illustrate the raw data and plot our estimates. In the figures, the solid line is our fit of Equation (1), while the hollow circles represent the raw average daily dependent variable for the days leading up to and after the tax change. For both the lines and circles, we remove the day-of-week, day-of-month, and firm location fixed effects. We also remove the estimates on the indicators for June 30 to July 5 from the fitted line.

[16] We cannot construct the optimal bandwidth in the style of Calonico, Cattaneo and Titiunik (2014) because of our day-of-week and day-of-month fixed effects (necessary due to strong cyclical patterns in the weight data) lead to perfect multicollinearity over small bandwidths.

4.1 Cultivators

We first focus on cultivators—the firms that plant and harvest raw flowers from the cannabis plant. Table 3 provides summary statistics for this group. Firms harvest plants, on average, 116 days after planting. Vertical integration is rampant, with over 95% of growers who were open before the tax change also having a license to process marijuana. The July 1 policy reform decreased their transaction tax from 25% to zero. However, because of widespread vertical integration, most cultivators did not owe this tax—taxes were not due if the cultivator sold to a vertically integrated processor. The high degree of vertically integrated transactions has also made the cultivator-to-processor transaction data unreliable.[17] We thus focus on other outcomes as vertically integrated cultivators may still respond to the decrease in the transaction tax between processors and retailers. These outcomes, which auditors routinely verify among growers, include counts of the number of new plantings, and the total quantities of marijuana harvested.

Table 4 provides point estimates for three outcomes of interest: the number of plantings, the number of harvests, and the average number of growth days for plants harvested. For both plantings and harvests, zeros are included in the analysis when appropriate. The left panel of Figure 1 illustrates the counts of the plants harvested around the July 1 tax change. The harvest rate increases by 7.5%, but this change is not statistically significant. We may, however, be underpowered to detect effects in the harvesting and planting data because these activities are performed infrequently and, when they are done, occasionally occur in large volumes. We examine the number of days from planting until harvesting, and find evidence that the number of days from planting until harvest falls by 6.9%. This suggests that cultivator-processors react to the processor tax decrease by speeding up the harvest of existing plants in order to get more to market in the medium-run.

We examine the sensitivity of our estimates in Table 5 and Figure 2. Our results are

[17]The administrative data does not contain meaningful prices for many of the transfers of plant material between cultivators and processors, and so we focus on quantity outcomes.

largely insensitive to bandwidth, allowing for an adjustment period from June 27 through July 5, as well as for polynomial order, interacting the polynomial with county indicators, and exclusion of day-of-month indicators.

These estimates suggest that in the short run, firms did not drastically alter their planting or harvesting in response to the tax shift. Given the vast amount of vertical integration and the fact that firms are producing at capacity due to strong demand, this is not entirely surprising—particularly given that the firms did not have a long period to adjust their production process since the law was passed mere days before it went into effect.

4.2 Processors

The tax reform also eliminated the gross receipts tax faced by processors—those firms that take raw marijuana flowers as inputs and transform them into usable marijuana. Before the change, the processors remitted the taxes after the transaction, implying that the retailers paid the list price, while processors received the after-tax price. Table 6 provides summary statistics for processors.

In Table 7, we provide point estimates of the effect of the tax change on the following potential mechanisms through which processors might have altered their behavior: prices (those charged to retailers), after-tax prices (those received by the processor after taxes), weight (total weight of marijuana sold in grams), sales (total number of sales), after-tax revenue, and THC levels.[18] The only significant findings concern prices—we find that average list prices fell by 8%. At the same time, the after-tax price (the price received by processors after taxes) increased by 21%.[19] We also find the quantity of sales transactions, total weight of marijuana sold, and THC levels were unchanged.[20]

[18]The number of observations differ for sales and prices, because days with no transactions still result in 0 sales or weight, while prices were missing on those days.

[19]The log change in the tax rate is 29% (i.e., log((1-.75)/1)=0.287).

[20]Within the "usable marijuana" product segment, THC levels are effectively fixed by the production process. However, processors may choose to differentially allocate plant material of varying quality to different production processes (including processes that result in products that fall outside of the "usable marijuana" category and thus outside of our analysis). Our null result here suggests that firms did not

Figure 3 provides graphical evidence of the after-tax price and weight responses we estimated in Columns (2) and (3) of Table 7. Because the after-tax price for processors increased dramatically following the tax change, it appears that the elimination of the 25% tax was a huge boon to processors, while it simultaneously decreased the marginal costs faced by retailers. At the same time, we find that the quantity of marijuana processed and sold to retail firms did not significantly change around the reform. This suggests firms did not anticipate the tax reform or its timing. The substantial shift in taxes should give firms incentives to hold back production and sales to retail firms until after the processor taxes are eliminated. Only the day before the tax change do we find that the quantity of marijuana substantially decreases (indeed, June 30 appears to be a true outlier, consistent with tax avoidance by processors—albeit extremely short-sighted avoidance).

Table 8 and Figure 4 detail several variations of our preferred specification for the significant after-tax price response observed in Column (2) and the insignificant change in weight observed in Column (3) of Table 7. The estimates are quite insensitive to bandwidth, allowing for an adjustment period from June 27 through July 5, as well as for polynomial order, interacting the polynomial with county indicators, and exclusion of day-of-month indicators.

In summary, our findings suggest that processors substantially benefited from the elimination of their transaction taxes with retailers. The benefits were shared: the average price paid by retailers fell by 7%, and the after-tax price received by processors rose by 21%. After we estimate the retail response to the reform in Section 4.3, we will discuss the tax invariance and incidence implications of these results. We find no significant increase in the quantity of marijuana supplied from processors to retailers, other than a one-day anticipation effect and a few days' adjustment period during which the inventory withheld the day before the tax was sold. This suggests that in the short run, the tax changes did not fundamentally alter the supply of marijuana. Taken together, these findings suggest that in the short-run window surrounding the tax change, supply is relatively inelastic, which is not surprising,

substitute in this way.

given the nature of the production process.

4.2.1 Vertical Integration of Cultivators and Processors

In this section, we provide the first empirical evidence for how vertical integration behavior adjusts to the elimination of a gross receipts tax. These estimates inform us about the potential deadweight loss associated with inefficient vertical integration stemming from the incentives created by a gross receipts tax. This setting is ideal for this analysis because we have data on each link in the supply chain and it is a closed system.[21]

We define vertical integration in this context at the retailer inventory-lot level. An inventory lot of marijuana is considered "vertically integrated" marijuana if it was cultivated and processed by the same firm. An inventory lot is considered "vertically dis-integrated" marijuana if it was cultivated and processed by different firms.[22]

Before the tax reform, cultivators and processors had a strong incentive to vertically integrate due to the 25% gross receipts tax that could be avoided. From accounts of industry participants, cultivators often wished they did not have to invest in the equipment to be processors, but could not afford to do otherwise. Moreover, even if a company was both a cultivator and a processor, it would have still have liked to sell some of its raw material to other processors or to have been able to purchase from other cultivators, but both options were often made too unattractive by the existence of the 25% tax. These accounts are borne out by the data, as 93% of all inventory lots sold by processors to retailers consisted of vertically integrated marijuana and only 5% of the total weight of processor sales were of vertically dis-integrated marijuana. However, vertical integration between cultivators and processors could be efficient, and so this alone does not inform us as to the magnitude of the inefficiency associated with the gross receipts tax system. The tax reform that came into effect on July 1, 2015 eliminated the incentive to vertically integrate for tax reasons. We do

[21] That is, the marijuana cultivators are not also, for example, selling to cigarette companies in an unobservable way

[22] Capacity constraints may lead a cultivator which is part of a vertically integrated firm to contract with a different processor which is part of a separate vertically integrated firm.

not expect that many firms would de-integrate, but new firms may be less likely to integrate and there may be more processors purchasing from cultivators with whom they were not vertically integrated.

To test this hypothesis, we examine the number of transactions of vertically integrated marijuana and the total weight of vertically dis-integrated marijuana. Neither of these measures should be expected to change immediately in response to the tax change because, on average, it takes six weeks after processors purchase raw material from cultivators before the resulting products are sold to retailers. Therefore, we adjust our analysis in this section relative to our previous processor analysis. First, we aggregate the analysis at the weekly level and examine a longer time span (45 weeks), and second, we include indicator variables for the six weeks following the tax change to differentiate the adjustment period from the long-run effect. In contrast to the previous analysis, we also do not drop new firms in this section because of the particular interest they present as being the firms most likely to choose to not be vertically integrated after the tax change.

Figure 5 illustrates the fraction of transactions that are vertically integrated, and the log weight of vertically dis-integrated marijuana. Point estimates are provided in Table 9. The fraction of vertically integrated sales falls by 3.7% after the adjustment period (Column 1), which is driven by a 42% long-run increase in non-vertically produced marijuana sold (Column 2). This provides evidence that the gross receipts tax discouraged otherwise efficient trades between cultivators and processors, thus creating deadweight loss. Column (3) of Table 9 shows that there was no significant increase in the price of these non-vertically integrated sales in the long run.

We further examine the potential inefficiencies created by the gross receipts tax by considering the behavior of each firm in its first week in business—in other words, by considering behavior closer to the firms' decision to vertically integrate—in Table 9, Columns (4) and (5).[23] The fraction of vertically integrated sales in the first week in business declines by 16%

[23]The responses for all firms are not driven solely by these new firms.

and the weight of vertically dis-integrated marijuana increases by 105% after the reform. Table 10 and Figure 6 illustrate the robustness of the results, with a variety of trend controls, changes to the adjustment periods, and variations in bandwidth.

In summary, we find that, in the long run, the volume of vertically dis-integrated transactions rose substantially, although vertically integrated transactions continue to dominate the market. Most of the estimates in this section are likely lower bounds on the true inefficiencies of a gross receipts tax as we are examining the elimination of a gross receipts tax. At the time of the reform, firms had already paid the fixed costs necessary to become both cultivators and processors, and therefore may not have chosen to dis-integrate afterwards, even if dis-integration is efficient in the absence of the gross receipts tax. The exception to this comes from firms in their first week in business. Not surprisingly, these estimates are more than twice as large.

4.3 Retail Market

In contrast to cultivators and processors, retailers are not allowed to vertically integrate with any other firms. The tax reform on retailers converted the retail tax from a gross receipts tax to an excise tax, and also switched the state sales tax from applying to the tax-inclusive price before the reform to the tax-exclusive price afterwards. The tax rate increased from 25% to 37%. Given that we can observe a portion of every retailer's input prices and those of their nearest competitors, we add some covariates to our regression analysis. We include the prices paid to processors, the prices of local[24] competitors, and an indicator for the presence of local competitors. These additions may be valuable particularly for estimates of price responses because, at the firm level, the observed price per gram could change not only due to a change in the product-level prices, but also due to a change in the composition of goods sold. Retailers also have significant weekend patterns in weight sold which may vary by firm based on their location and hours of operation. We include interactions between each

[24]We define 'local' as 'within 10 miles.'

firm indicator and indicators for Friday, Saturday, and Sunday which increases precision. Summary statistics for the retail market are provided in Table 11.

Point estimates across a variety of outcomes are detailed in Table 12. We estimate the response of tax-exclusive and -inclusive prices, weight, the number of sales, tax-exclusive and -inclusive revenue, and THC content. Tax-exclusive prices—the prices retailers received—fell by 4.8% while tax-inclusive prices—the prices paid by consumers—rose by 2.1%.[25] We find no significant response in weight or sales to this increase in the prices consumers paid. We do find a decline in THC content which, when combined with the null THC result in the processor analysis, suggests that consumers substituted toward lower quality marijuana in response to the increase in tax-inclusive prices and our processor price controls are not able to fully control for these substitution patterns. This biases our price estimates downwards: firms likely increased sticker prices by a larger amount than a regression using purchases would detect.

We address this bias by collapsing our data by inventory lot-day instead of by retail firm-day and including inventory-lot fixed effects rather than retail-location fixed effects.[26] Because inventory lots are unique to individual retailers and consist of plausibly identical goods, quality (including retailer quality) within a lot is fixed. In addition to controlling for bias created by substitution, we also anticipate that our estimates will be substantially more efficient. We do not estimate the weight response at the inventory-lot level as a substitution from one gram of high quality marijuana to one gram of low quality marijuana does not change the total weight sold. Indeed, analyzing the weight response at the inventory-lot level could be misleading since we estimate the *log* change, and increases and decreases of one unit each are not equivalent unless the sales volume of the two inventory lots started

[25]The excise tax increased from 25% to 37% and simultaneously the sales tax switched from being assessed on the tax-inclusive price to being assessed on the tax-exclusive price. Together, these amount to a 7% increase in the tax rate ($1+\tau$) as measured by the log change in $1+\tau$ for a firm with the average sales tax rate of 8.7% ($= \log((1 + .37 + .087)/(1.25 * 1.087))$.

[26]We restrict this analysis to lots that sell at least once in the month prior to the tax change. Our estimates are identified from lots that sell both before and after the tax change regardless; including this restriction allows us to more accurately estimate the coefficients for the lots that provide our identifying variation.

out the same.[27] Our preferred estimates of the increase in the tax-inclusive price using the inventory lot-day level data is reported in Table 13, Column (1). We find that tax-inclusive prices increased by 2.3% and that this estimate is significant at the one-percent level. The magnitude of the inventory-lot and firm level estimates are quite similar, suggesting that processor prices were, in fact, doing a fairly good job of capturing the substitution patterns we were concerned about at the firm level.

Combining our preferred estimates on the change in the tax-inclusive price (estimated at the inventory lot-day level) and weight (estimated at the retailer-day level) reported in Table 13, Column (1), the implied average market-level elasticity of demand is -0.43. Figure 7 illustrates these shifts in tax-inclusive price and total weight surrounding the tax reform. For the regression, which we report in Table 13, we switch to a linear polynomial in date sold with an interaction between the tax change and this linear polynomial. The estimates are not sensitive to this change in polynomial (see Figure 9) and this allows us to conserve degrees of freedom and to assess whether the tax change also affected the steady decline in prices or the increase in weight over time. This short-run price elasticity of demand suggests that Washington is on the part of the Laffer curve where higher taxes increase revenue, and this is consistent with the positive effect on tax-inclusive revenue in Table 12, Column (6). The decrease in tax-*exclusive* prices, however, leads to a significant decrease in revenue received by retailers. The long-run elasticity of demand appears higher. There is a significant negative slope on the weight interaction term, which means that the weight is continuing to decline over time. In fact, the implied elasticity of demand is about negative one within two weeks of the tax change. Although we do not have the power to detect this in our estimation, visual inspection of the raw data in Figure 7 suggests the decline in the growth rate of marijuana demand slows down several weeks after the tax change, suggesting a cut-off for the medium-run elasticity of demand, holding all else constant. This larger medium-run elasticity of demand suggests that Washington is near the peak of the Laffer curve.

[27]The estimates of the weight response when the data are collapsed by inventory lot are qualitatively similar to such estimates when the data is collapsed by firm.

We consider the sensitivity of our main price and weight estimates in Table 13, Columns (2) to (7) and in Figures 8 and 9. Our estimates are insensitive to bandwidth, allowing for an adjustment period of June 27 through July 5, for decreasing the adjustment period to June 30 through July 1, and for interacting the polynomial with county indicators. We also show that the weight estimates are insensitive to dropping the inclusion of processor price covariates, which is not surprising given that the weight estimates should not be influenced by substitution decisions on the part of consumers, as discussed above. We report the baseline price estimate again in this column because we never control for processor prices in the inventory lot-day analysis (these covariates would be collinear with the inventory-lot fixed effects). The weight estimates are sensitive to the exclusion of the day-of-month indicators in Column (4), which is not surprising as there are systematic fluctuations in demand within a month. There is no noticeable effect on the price estimates when we drop the interactions of Friday, Saturday, and Sunday with each retail firm location in Column (5). The weight estimates get slightly smaller and less precise.

With estimates of the changes in prices for both the processor and retailer markets in hand, we summarize the components of the tax-inclusive price charged by the average firm for a gram of marijuana in Figure 10. The left and right bars, respectively, present the components of the retail price before and after the tax change. Starting with the average price paid for a gram of marijuana, we consider how much goes to processor and retail taxes as well as how much the firms keep. Before the tax change, all prices and taxes (in dollars) are based on the average prices across all firms the month prior to the tax change. The post-reform numbers are calculated by applying our estimated price responses to the pre-reform averages. This holds the composition of the market constant and eliminates any secular trends in prices as well as changes in the consumption patterns of consumers.

The tax changes that were implemented on July 1, 2015 were intended to be approximately revenue neutral, and, in practice, they were not too far off from this. We calculate that the average excise taxes collected on a gram of marijuana before the tax change were

$3.46 ($1.03 collected and paid by the processor and $2.43 paid by the retailer) and that the average taxes paid on a gram of marijuana after the tax change were $3.42 (all collected and paid by the retailer).[28] When we add sales taxes into this mix, which switched from being applied on the tax-inclusive price before the tax change to tax-exclusive price after the tax change, the total excise and sales taxes collected before the tax change were $4.51 per gram of marijuana and were $4.23 afterwards, a decrease of 28 cents. Total retail taxes (leaving the processor tax change out of the mix) went up by 75 cents.

Processors paid an average of $1.03 in taxes for each gram sold prior to the change (and nothing afterwards). The tax reform eliminated 28 cents of taxes overall, but the rest were shifted from the processor to the retailer. The tax-invariance folk theorem suggests that, in light of the fact that incidence is fixed and already accounted for in pre-reform prices, this shift in taxes from the processor to the retailer should be reflected in the processor price falling by the full amount of the tax shifting, 75 cents. We find, however, that processors only lower their price by an average of 26 cents per gram, suggesting tax invariance has been violated in this setting.

We now consider the incidence implications of our results. If we were only considering a retail tax change (i.e., the excise tax rate increased from 25 to 37% and the sales tax switched from being levied on the tax-inclusive to the tax-exclusive price), the incidence calculations would be straightforward. The retail tax rate increased by 7%[29], and prices faced by consumers went up by 2.33%, so tax pass-through to the consumer would be 33%.

However, the processor tax (and cultivator tax for unintegrated firms) also changed, which led to a decline in processor prices (i.e., retailers' marginal costs went down). We expect that marginal cost changes are passed through to consumers in the same way as tax changes, so our calculation of consumer incidence that ignores this processor price change is

[28] We ignore the tax applied to cultivators in this discussion because over 90% of the market was vertically integrated, and thus did not have to pay the tax. Our price data for the remaining cultivators is in any case not of high-enough quality to consider this part of the tax change directly.

[29] The tax changes amount to a 7% increase in the tax rate $(1+\tau)$ as measured by the log change in $1+\tau$ for a firm with the average sales tax rate, 8.7% $(= \log((1 + .37 + .087)/(1.25 * 1.087))$.

biased downwards. In fact, if the tax change was revenue neutral and tax-invariance held, we would expect no change in retail prices, not because there was no pass-through, but because there was no change in taxes to be passed-through. Figure 10 shows average processor prices paid by retailers fell by 26 cents per gram. This amounts to a 1.94% decrease in marginal costs paid by the retailer from the equilibrium tax-inclusive retail price, \$13.49.[30]. Hence, the net change in taxes and marginal costs is 5.06%, so consumers bear 44% of the tax burden (=.0233/.0506).[31]

5 Conclusion

Recreational marijuana will soon be legal for 21% of the United States' population. Given the broad society-wide shift in support for legalization, driven in part by the desire to increase state revenues, it is crucial to accurately gauge how industry participants, and, ultimately, revenue, will respond to alternative policies. Internationally, countries that have chosen to legalize marijuana, such as Canada and Uruguay, have had robust debates about regulatory and taxation design in the face of black market availability.

We contribute to these debates by analyzing a unique natural experiment in Washington, where the tax regime was changed from a three-tiered structure of gross receipts taxes throughout the supply chain to a single excise tax at retail. We use novel data on the universe of marijuana production and sales within the state to understand how cultivators, processors, retailers, and consumers responded to this reform, and how their responses affected the market prices at the various stages of production and distribution.

We find that gross receipts taxes increased vertical integration, which decreased the efficiency of the market, given that non-vertically integrated transactions increased substantially, and that newly entering firms were far less likely to vertically integrate after the tax system moved away from a gross receipts tax. This suggests that the deadweight loss of

[30]We calculate this as the log change, $0.0194 = \log((13.49 - .26)/13.49)$.

[31]We examine how incidence may be affected by the federal income tax in Appendix A.

these gross receipts taxes may be larger than alternative revenue-neutral taxes. This finding is particularly relevant given the diverse regulations present across the United States. Colorado requires vertical integration (retail firms must produce at least 80% of the marijuana they sell) and Washington allows some types of vertical integration (between cultivators and processors) while forbidding integration between retailers with firms involved in marijuana production. Our findings suggest that requiring vertical integration decreases efficiency.

Though seven of the eight states with legal marijuana have a sales or excise tax, several states have also imposed cultivation taxes. Our findings suggest that cultivation taxes can directly affect incentives to vertically integrate. Additionally, Alaska and California's cultivation taxes are based on weight alone and Colorado's wholesale tax is based on the average price-per-gram of marijuana sold in the state. As the amount taxed in those states does not change in response to the price of the wholesale marijuana sold, these taxes are relatively higher for cheaper marijuana, which tends to be of lower quality and potency. The tax regimes in Alaska, California, and Colorado may thus incentivize suppliers to produce higher priced and higher potency marijuana.

While the previous literature has studied the elasticity of demand for marijuana within the black market, we are the first to focus on the legal market. Some recent studies, notably Jacobi and Sovinksy (2016), have estimated the elasticity of marijuana to be -0.2. Our findings suggest that demand is inelastic (-0.4) immediately after the price increase driven by the exogenous tax reform. However, we find that the quantity of marijuana purchased kinks downward from its prior upward trend. This implies that the medium-run response to a price increase is elastic, which in turn suggests that Washington is close to the peak of the Laffer curve. The tax rates that were passed in California (15%), Maine (10%), Massachusetts (10.75%), and Nevada (15%), are considerably lower than Washington's rate. If supply and demand characteristics are similar[32] once these newer markets mature (our data cover a period approximately one year after the initial legalization came into effect),

[32]While supply characteristics are likely similar, given the prevalence of cultivators using indoor technology, demand may differ due to differing incomes and tourist activity.

our results suggest that significant state revenue may be left on the table in these other states.

Though tax revenue has historically been one of the many arguments in favor of legalizing marijuana, evaluating the impact of marijuana policy (and constructing optimal policy) in a broader social sense requires additional considerations. For one, the public health externalities of marijuana consumption are not well established; nor is the relationship between legal marijuana consumption and the consumption of other "sin" goods, such as alcohol or tobacco. If it is indeed true, as many advocates claim, that marijuana consumption is 'better' in a public health sense than alcohol or tobacco consumption, the optimal regulation of marijuana should be designed to take into account responses in these other markets as well. We leave such analysis to future authors.

References

[1] Anderson, D. M., B. Hansen, and D. I. Rees. (2013). "Medical Marijuana Laws, Traffic Fatalities, and Alcohol Consumption." *Journal of Law and Economics*. 56.2 : 333-369.

[2] Anderson, D.M., B. Hansen and D.I. Rees. (2015). "Medical Marijuana Laws and Teenage Marijuana Use." *American Law and Economics Review*. 17 (2): 495-528.

[3] Barnett, P. G., Keeler, T. E., & Hu, T. W. (1995). "Oligopoly Structure and the Incidence of Cigarette Excise Taxes." *Journal of Public Economics*, 57(3), 457-470.

[4] Brockmeyer, A. & Hernandez, M. (2016). Taxation, Information, and Withholding: Evidence from Costa Rica (No. 7600). World Bank Policy Research.

[5] Calonico, S., Cattaneo, M. D., & Titiunik, R. (2014). Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs. *Econometrica*, 82(6), 2295-2326.

[6] Chetty, R., Looney, A. and Kroft, K (2009). "Salience and Taxation: Theory and Evidence." *American Economic Review*, 99(4), 1145-1177.

[7] Chaloupka, F. J., Grossman, M., & Tauras, J. A. (1999). "The Demand for Cocaine and Marijuana by Youth. In The Economic Analysis of Substance Use and Abuse: An Integration of Econometrics and Behavioral Economic Research"(pp. 133-156). *University of Chicago Press*.

[8] Clements, K. W., & Zhao, X. (2009). Economics and Marijuana: Consumption, Pricing and Legalization. Cambridge University Press.

[9] Cole, J. (2013). Memorandum for all United States attorneys. accessed from https://www.justice.gov/iso/opa/resources/305201382913275685746 7.pdf on 07/20/2016

[10] Dobkin, C., & Nicosia, N. (2009). "The War on Drugs: Methamphetamine, Public Health, and Crime." *The American Economic Review*, 99(1), 324-349

[11] Donohue III, J. J. (2013). "Drug Prohibition and Its Alternatives. Lessons from the Economics of Crime: What Reduces Offending? "*Lessons from the Economics of Crime'in Lessons from the Economics of Crime*, edited by PJ Cook, SJ Machin, O. Marie, and G. Mastrobuoni, MIT Press.

[12] Farrelly, M. C., Bray, J. W., Zarkin, G. A., Wendling, B. W., & Pacula, R. L. (1999). The Effects of Prices and Policies on the Demand for Marijuana: Evidence from the National Household Surveys on Drug Abuse (No. w6940). National Bureau of Economic Research.

[13] Hahn, J., Todd, P., & Van der Klaauw, W. (2001). Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design. *Econometrica*, 69(1), 201-209.

[14] Hall, W., & Lynskey, M. (2016). "Evaluating the Public Health Impacts of Legalizing Recreational Cannabis Use in the United States. "*Addiction*, 111:1764-1773.

[15] Hansen, B., Miller, K., & Weber, C. (2017) "How Extensive is Cross-State Diversion of Recreational Marijuana? "*Working paper*.

[16] Harding, M., Leibtag, E., & Lovenheim, M. F. (2012). "The Heterogeneous Geographic and Socioeconomic Incidence of Cigarette Taxes: Evidence from Nielsen Homescan Data. "*American Economic Journal: Economic Policy*, 4(4), 169-198.

[17] Hausman, C. & Rapson, D. S. (2017). Regression Discontinuity in Time: Considerations for Empirical Applications. Technical report, National Bureau of Economic Research.

[18] Jacobi, L. & M. Sovinsky. (2016). "Marijuana on Main Street? Estimating Demand in Markets with Limited Access."*American Economic Review*. 106(8): 2009-45.

[19] Kaeding, N. (2017) "The Return of Gross Receipts Taxes."*www.taxfoundation.org*. Tax Foundation, 28 March, 2017.

[20] Kenkel, D. S. (2005). "Are Alcohol Tax Hikes Fully Passed through to Prices? Evidence from Alaska. "*The American Economic Review*, 95(2), 273-277.

[21] Kopczuk, W., Marion, J., Muehlegger, E., & Slemrod, J. (2016). "Does Tax-Collection Invariance Hold? Evasion and the Pass-through of State Diesel Taxes. "*American Economic Journal: Economic Policy*, 8(2), 1-36.

[22] Kopczuk, W. & Slemrod, J. (2006). "Putting Firms into Optimal Tax Theory. "*American Economic Review*, 96(2), 130-134.

[23] Laffer, A. (2004). "The Laffer Curve: Past, Present, and Future."

[24] La Corte, R. (2015) "Gov. Inslee signs recreational marijuana reform bill. " *ap.org*. Associated Press, 30 Jun. 2015.

[25] McCrary, J. (2008). Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test. *Journal of Econometrics*, 142(2), 698-714.

[26] McLure, C., (2005). "Why Ohio Should Not Introduce a Gross Receipts Tax Testimony on the Proposed Commercial Activity Tax." *State Tax Notes*, Apr. 18, 2005, p. 213.

[27] Miron, J., & Zwiebel, J. (1991). "Alcohol Consumption During Prohibition." *American Economic Review*, 81(2), 242-247.

[28] Miron, J. A., & Zwiebel, J. (1995). "The Economic Case Against Drug Prohibition. *The Journal of Economic Perspectives* 9(4), 175-192.

[29] Moffitt, Thomas D. (2015). "Memorandum: Section 164, Section 280E, and the State of Washington Marijuana Excise Tax." accessed from `https://www.irs.gov/pub/irs-wd/201531016.pdf` on 08/01/2017

[30] Motel, S. (2015). Six Facts about Marjuana. accessed from http://www.pewresearch.org/fact-tank/2015/04/14/6-facts-about-marijuana/ on 10/22/2016.

[31] Ogden, D. (2009). Memorandum for selected United States Attorneys. accessed from https://www.justice.gov/archives/opa/blog/memorandum-selected-united-state-attorneys-investigations-and-prosecutions-states on 07/20/2017

[32] Pogue, T. (2007). "The Gross Receipts Tax: A New Approach to Business Taxation," *National Tax Journal*, 60(4), 799-819.

[33] Poterba, J. (1996). "Retail Price Reactions to Changes in State and Local Taxes. *National Tax Journal*, 49, pp. 165-176.

[34] Slemrod, J. (2008). "Does It Matter Who Writes the Check to the Government? The Economics of Tax Remittance. " *National Tax Journal*, 61(2), pp. 251-75.

[35] Testa, W. & R. Mattoon (2007). "Is There a Role for Gross Receipts Taxation?" *National Tax Journal*, 60(4), 821-840.

[36] Williams, J., Van Ours, J. C., & M. Grossman. (2011). "Why Do Some People Want to Legalize Cannabis Use?" (No. w16795). *National Bureau of Economic Research.*

Tables

Table 1: Current marijuana taxes by state

State	Tax rate	Notes
Alaska	None	Cultivation tax of $50/oz on dried flowers and $15/oz on for other plant material.
California	15%	Cultivation tax of $9.25/oz on dried flowers and $2.75/oz on dried leaves.
Colorado	10%	Additional 15% tax applied at wholesale based on average market rate in the state.
Maine	10%	
Massachusetts	10.75%	Localities may impose additional 3% tax.
Nevada	15%	Tax applied at wholesale.
Oregon	17%	Localities may impose additional 3% tax.
Washington	37%	

Note: All taxes applied to retail sales unless otherwise noted. Washington initially set a tax rate of 25% both at retail and for transfers within the supply chain. On July 1, 2015, Washington switched to a single 37% tax at retail.

Table 2: Summary of Washington House Bill 2136 legislative history

Date	Activity
	Regular Session
February 17	Bill introduced
February - March	Committee work results in significant changes to bill
April 10	Passed by House 67-28, referred to Senate
April 24	Regular Session ends without vote
	First Special Session
April 29	House passes 70-25, referred to Senate
May 1	Referred to Senate committee
May 28	First Special Session ends without vote
	Second Special Session
May 29	Reintroduced in House, referred to committee
June 26	Passed by House 59-38
June 27	Passed by Senate 36-7
June 30	Signed by Governor
July 1	**Law takes effect and tax rates change**

Source: `http://app.leg.wa.gov/billsummary?BillNumber=2136&Year=2015`

Table 3: Cultivator Summary Statistics

	Obs.	Mean	Std. Dev.	Mean>0	Min.	Max.
Plantings	45,401	8.94	66.69	93.59	0	3,646
Harvesting	30,477	7.78	66.70	63.39	0	4,153
Plant to Harvest Days	3,739	116.20	40.61	116.20	2	292
Firm has Processor License?	250	0.96	0.19	0.96	0	1

Plantings and harvesting are averages over cultivator location-day for all days after which a cultivator location had first planted (or harvested) in the two months before and after the tax reform. These variables are zero if no planting or harvesting occurred on that day. Plant-to-harvest days is an average over cultivator location-day of the number of days from planting to harvesting for every cultivator location-day in which some plants were harvested. Whether a firm has a processor license is a one-time calculation of the fraction of cultivators that have a processor license. All variables are subject to the data cleaning described in our Online Data Cleaning Appendix: `http://www.keatonmiller.org/s/data-cleaning-appendix.pdf`.

Table 4: Cultivator Response to Tax Change

	(1) Plantings	(2) Harvests	(3) Days-to-Harvest
Tax Change	−0.005	0.075	−0.069*
	(0.035)	(0.051)	(0.036)
Observations	45,401	30,477	3,739
R-squared	0.106	0.104	0.569
Cultivators	374	250	250

The following variables are included, but not reported in these regressions: day indicator variables for June 30 - July 5, day of the week and day of the month indicator variables, a 3rd-order polynomial in planting or harvesting date, and cultivator location fixed effects. All dependent variables are in logs (or the log of 1 plus the variable if there are zeros); the dependent variable is listed directly below the column number. Standard errors clustered by cultivator are in parentheses.

Table 5: Robustness Checks for Cultivator Response to Tax Change

	(1)	(2)	(3)	(4)	(5)	(6)
Harvests	0.075	0.068	0.057	0.087*	0.084	0.074
	(0.051)	(0.057)	(0.045)	(0.048)	(0.066)	(0.051)
Days-to-Harvest	−0.069*	−0.085**	−0.058*	−0.048	0.003	−0.074**
	(0.036)	(0.037)	(0.031)	(0.032)	(0.050)	(0.036)
Harvests Observations	30,477	30,477	30,477	30,477	30,477	30,477
Days-to-Harvest Observations	3,739	3,739	3,739	3,739	3,739	3,739
Cultivators	250	250	250	250	250	250
Indicators for June 27-29	No	Yes	No	No	No	No
Indicators for July 2-5	Yes	Yes	No	Yes	Yes	Yes
Day of Month Indicators	Yes	Yes	Yes	No	Yes	Yes
Polynomial Order	3rd	3rd	3rd	3rd	5th	3rd
County x Polynomial Trend	No	No	No	No	No	Yes

The top row reports an estimate for the dependent variable harvestings. The next row reports an estimate for the dependent variable days-to-harvest. Dependent variables are logged (or the log of 1 plus the variable if there are zeros). The following variables are included, but not reported in these regressions: day indicator variables for June 30 - July 5, day of the week and day of the month indicator variables, a polynomial in harvesting date, and cultivator location fixed effects. Standard errors clustered by cultivator are in parentheses.

Table 6: Processor Summary Statistics

	Obs.	**Mean**	**Std. Dev.**	**Mean>0**	**Min.**	**Max.**
Price per Gram	6,522	3.65	1.12	3.65	0.01	10.62
Weight (in grams)	27,424	358.55	1,179.40	1,507.66	0	26,558.50
Number of Sales	27,424	3.34	10.47	14.06	0	270
Firm Revenue	27,424	1319.62	4274.30	5,548.81	0	100,622.17
Vertically Integrated?	6,514	0.93	0.22	0.93	0	1

Weight, number of sales, and firm revenue are averages over processor location-day for all days after which a processor location had first made a sale to a retailer in the two months before and after the tax reform. These variables are zero if no sales occurred on that day. Price per gram and the fraction of vertically integrated firms is an average over cultivator location-day of these variables for every cultivator location-day in which a sale took place. All variables are subject to the data cleaning described in our Online Data Cleaning Appendix: `http://www.keatonmiller.org/s/data-cleaning-appendix.pdf`.

Table 7: Processor Response to Tax Change

	(1) Price	(2) AT Price	(3) Weight	(4) Sales	(5) Rev.	(6) AT Rev.	(7) THC
Tax Change	−0.076**	0.212***	0.009	0.009	−0.009	0.054	−0.036
	(0.034)	(0.034)	(0.085)	(0.030)	(0.102)	(0.100)	(0.093)
Observations	6,522	6,522	27,424	27,424	27,424	27,424	6,399
R-squared	0.269	0.301	0.233	0.271	0.232	0.233	0.832
Processors	225	225	225	225	225	225	224

The following variables are included, but not reported in these regressions: day indicator variables for June 30 - July 5, day of the week and day of the month indicator variables, a polynomial in processor sale date, and processor location fixed effects. All dependent variables are in logs (or the log of 1 plus the variable if there are zeros); the dependent variable is listed directly below the column number. AT stands for after-tax. Standard errors clustered by processor are in parentheses.

Table 8: Robustness Checks for Processor Response to Tax Change

	(1)	(2)	(3)	(4)	(5)	(6)
After-Tax Price	0.212***	0.209***	0.221***	0.238***	0.189***	0.213***
	(0.034)	(0.035)	(0.030)	(0.032)	(0.054)	(0.034)
Weight	0.009	−0.018	0.047	0.014	0.129	0.006
	(0.085)	(0.095)	(0.078)	(0.079)	(0.114)	(0.085)
After-Tax Price Observations	6,522	6,522	6,522	6,522	6,522	6,522
Weight Observations	27,424	27,424	27,424	27,424	27,424	27,424
Processors	225	225	225	225	225	225
Indicators for June 27-29	No	Yes	No	No	No	No
Indicators for July 2-5	Yes	Yes	No	Yes	Yes	Yes
Polynomial Order	3rd	3rd	3rd	3rd	5th	3rd
County x Polynomial	No	No	No	No	No	Yes

The top row reports an estimate for the dependent variable tax-inclusive price per gram. The next row reports an estimate for the dependent variable after-tax revenue. Dependent variables are logged (or the log of 1 plus the variable if there are zeros). The following variables are included, but not reported in these regressions: day indicator variables for June 30 - July 5, day of the week and day of the month indicator variables, a 3rd-order polynomial in processor sale date, and processor location fixed effects. Standard errors clustered by processor are in parentheses.

Table 9: Vertical Integration Response to Tax Change

	(1) Vertical	(2) NVWeight	(3) NVPrice	(4) Vertical	(5) NVWeight
Tax Change	−0.037**	0.418***	0.047	−0.162**	1.046**
	(0.018)	(0.147)	(0.036)	(0.081)	(0.517)
One Week Post	0.041***	−0.378***	−0.079**	0.199***	−1.298***
	(0.015)	(0.138)	(0.037)	(0.056)	(0.359)
Two Weeks Post	0.037**	−0.342***	−0.007	0.195***	−1.271***
	(0.016)	(0.122)	(0.053)	(0.054)	(0.347)
Three Weeks Post	0.031**	−0.272***	−0.050	−0.059	0.083
	(0.013)	(0.104)	(0.036)	(0.225)	(1.208)
Four Weeks Post	0.028**	−0.192*	0.018	0.186***	−1.215***
	(0.011)	(0.108)	(0.039)	(0.050)	(0.322)
Five Weeks Post	0.012	−0.163*	−0.063**	0.182***	−1.188***
	(0.010)	(0.094)	(0.030)	(0.048)	(0.310)
Six Weeks Post	0.017	−0.189*	−0.048	0.178***	−1.160***
	(0.011)	(0.097)	(0.039)	(0.046)	(0.299)
Observations	16,422	23,940	2,549	523	523
R-squared	0.652	0.497	0.214	0.015	0.013
First Week Only	No	No	No	Yes	Yes
Weeks Pre-Post	45	45	45	45	45
Polynomial Order	3rd	3rd	3rd	3rd	3rd
Processor Locations	525	525	155	510	510

Processor location fixed effects (Columns (1) to (3) only) and a polynomial in the processor sale week are included, but not reported. NVWeight stands for log of 1 plus non-vertical weight. NVPrice stands for the ratio of the average vertical wholesale price over the average price for each processor location. Standard errors clustered by processor location are in parentheses for Columns (1) to (3) and heteroskedastic robust standard errors are in parentheses for Columns (4) and (5).

Table 10: Robustness Checks for Vertical Integration Response to Tax Change

	(1)	(2)	(3)	(4)	(5)	(6)
Fraction Vertically Integrated	−0.037**	−0.038**	−0.031*	−0.040**	−0.039*	−0.034**
	(0.018)	(0.019)	(0.017)	(0.019)	(0.021)	(0.017)
Non-Vertical Weight	0.418***	0.427***	0.351***	0.406***	0.488***	0.407***
	(0.147)	(0.154)	(0.133)	(0.156)	(0.159)	(0.145)
Fraction Vertically Integrated Observations	23,940	23,940	23,940	23,940	23,940	23,940
Non-Vertical Weight Observations	16,422	16,422	16,422	16,422	16,422	16,422
Processors	525	525	525	525	525	525
Number of Post-Reform Week Indicators	6	6	4	8	6	6
Number of Pre-Reform Week Indicators	0	1	0	0	0	0
Polynomial Order	3rd	3rd	3rd	3rd	5th	3rd
County x Polynomial	No	No	No	No	No	Yes

The top row reports an estimate for the dependent variable fraction of vertically integrated transactions by firm. The next row reports an estimate for the dependent variable log of 1 plus non-vertical weight. Processor location fixed effects and a polynomial in the week are included but not reported. Standard errors clustered by processor are in parentheses.

Table 11: Retail Summary Statistics

	Obs.	**Mean**	**Std. Dev.**	**Median**	**Min.**	**Max.**
Price per Gram	15,893	9.39	1.5	9.30	0.31	16.60
Weight (in grams)	15,893	511.89	381.41	464.44	1	7,434.61
Number of Sales	15,893	239.07	215.26	181	1	2,017
Firm Revenue	15,893	4479.69	4,043.55	3,335.62	4	36,854.39
Days from Wholesale	15,893	21.23	15.82	17.2	-9	189
THC Potency	15,893	19.66	1.45	19.73	9.05	32.76
CBD Potency	15,893	0.40	0.26	0.34	0	3.90
Number of Competitors	15,893	5.28	4	5	1	18

These variables are averages over retail location-day for all days in which sales occurred after which a retailer location had first made a sale to a consumer in the two months before and after the tax reform. There are no zeros included for days in which a firm did not sell because these generally only happen when they are closed on Sundays or Mondays, for example. All variables are subject to the data cleaning described in our Online Data Cleaning Appendix: `http://www.keatonmiller.org/s/data-cleaning-appendix.pdf`.

Table 12: Retail Response to Tax Change

	(1) Price	(2) TI Price	(3) Weight	(4) Sales	(5) Rev.	(6) TI Rev.	(7) THC
Tax Change	−0.048***	0.021*	0.007	0.022	−0.036	0.034	−0.018***
	(0.012)	(0.012)	(0.027)	(0.029)	(0.030)	(0.030)	(0.006)
Observations	15,893	15,893	15,893	15,893	15,893	15,893	15,893
Retailers	135	135	135	135	135	135	135
R-squared	0.794	0.778	0.885	0.897	0.887	0.888	0.494
Polynomial Order	5th	5th	5th	5th	5th	5th	5th

The following variables are included, but not reported in these regressions: log processor price, whether any competitors, log competitors' processor price, day indicator variables for June 30 - July 5, day of the week indicator variables (the Friday, Saturday, and Sunday indicator variables are interacted with each retail location fixed effect), day of the month indicator variables, a polynomial in retail sale date, and retail location fixed effects. All dependent variables are in logs; the dependent variable is listed directly below the column number. TI stands for tax-inclusive. Standard errors clustered by retailer are in parentheses.

Table 13: Robustness Checks for Retail Response to Tax Change

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Inventory Lot TI Price	0.0233***	0.0228***	0.0220***	0.0237***	0.0233***	0.0233***	0.0229***	0.0241***
	(0.0066)	(0.0068)	(0.0062)	(0.0064)	(0.0066)	(0.0066)	(0.0066)	(0.0068)
Inventory Lot TI Price Interaction	−0.0001	−0.0001	−0.0000	−0.0001	−0.0001	−0.0001	−0.0001	
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	
Weight	−0.0095	−0.0038	−0.0079	0.0135	−0.0085	−0.0092	−0.0096	−0.0072
	(0.0228)	(0.0229)	(0.0207)	(0.0215)	(0.0220)	(0.0222)	(0.0228)	(0.0228)
Weight Interaction	−0.0010*	−0.0008	−0.0010*	−0.0005	−0.0009	−0.0009	−0.0010*	
	(0.0006)	(0.0006)	(0.0005)	(0.0006)	(0.0006)	(0.0006)	(0.0006)	
Implied Elasticity	−0.4260	−0.1666	−0.3591	0.5696	−0.3812	−0.3948	−0.4192	−0.2988
Observations	15,893	15,893	15,893	15,893	15,893	15,893	15,893	15,893
Inventory Lot Observations	567,425	567,425	567,425	567,425	567,425	567,425	567,425	567,425
Retailers	135	135	135	135	135	135	135	135
Indicators for June 27-29	No	Yes	No	No	No	No	No	No
Indicators for July 2-5	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes
Day of Month Indicators	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes
Weekend Days x Retail Location	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes
Processor Prices Included?	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes
# Competitors Included?	No	No	No	No	No	No	Yes	No
County x Polynomial	No	No	No	No	No	No	No	Yes

The top two rows report estimates for the dependent variable log of the tax-inclusive price per gram where the data are aggregated by retail location-inventory lot-day instead of retail location-day. The first row reports the estimate on the *Tax Change* indicator variable. The second row reports the estimate on the interaction between the running variable and *Tax Change*. The next two rows report the same estimates for the dependent variable log weight where the data are aggregated by retail location-day. The elasticity of demand implied by these estimates is calculated by dividing the weight estimates in the third row by the price estimates in the first row. The following variables are included, but not reported in these regressions unless otherwise specified: log processor price, log competitors' processor price, whether any competitors (these first three covariates are included only for the third and fourth rows), day indicator variables for June 30 - July 5, , day of the week indicator variables (the Friday, Saturday, and Sunday indicator variables are interacted with each retail location fixed effect), day of the month indicator variables, and retail location fixed effects. All dependent variables are in logs. Standard errors clustered by retailer are in parentheses.

Figures

Figure 1: Cultivator Harvests and Days-to-Harvest

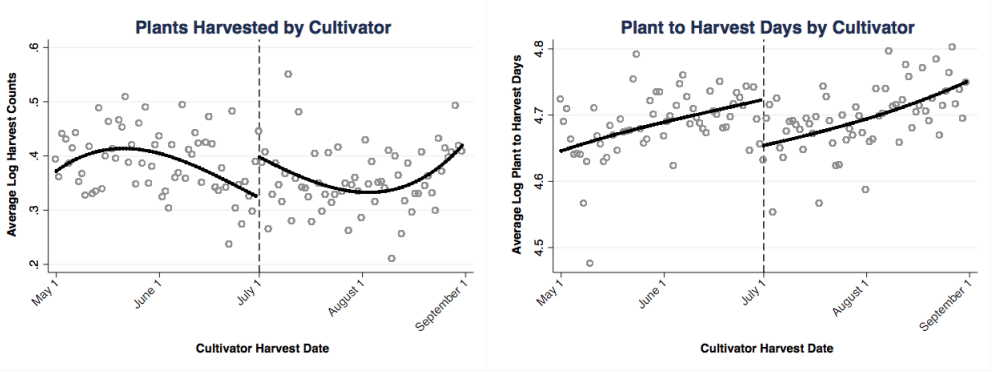

These figures are based on the estimates in Column (1) of Table 5. The solid line plots the estimated response to the tax change plus the polynomial. The scatterplot is the corresponding daily average of the dependent variable with estimates of the day-of-week, day-of-month, and cultivator fixed effects removed. The vertical dashed line marks the day of the tax change, July 1, 2015.

Figure 2: Cultivator Harvests and Days-to-Harvest Bandwidth Sensitivity

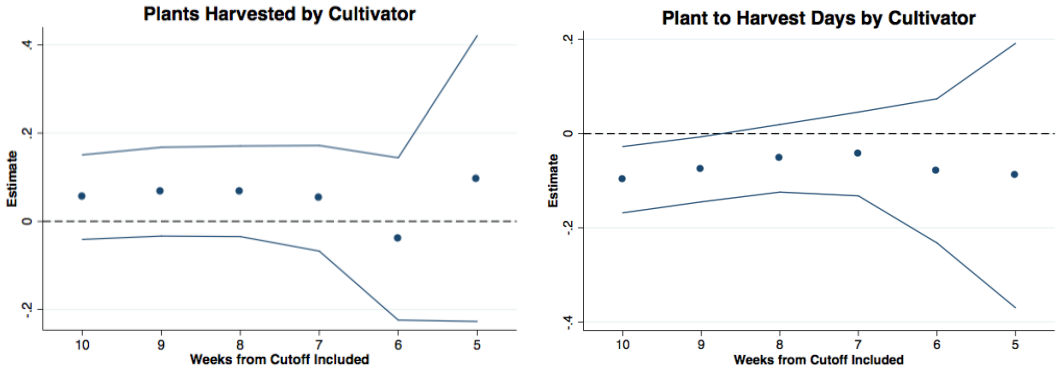

These figures consider the sensitivity of the estimates in Column (1) of Table 5 to the number of weeks of data we include on either side of the tax change. The dots mark the estimates for each bandwidth choice and the lines mark the 95% confidence intervals around these estimates.

Figure 3: Processor After-Tax Prices and Revenue

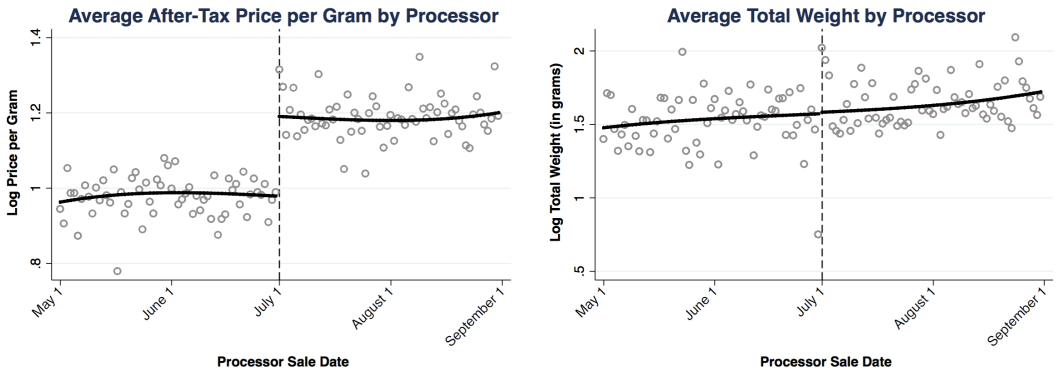

These figures are based on the estimates in Column (1) of Table 8. The solid line plots the estimated response to the tax change plus the polynomial. The scatterplot is the corresponding daily average of the dependent variable with estimates of the processor fixed effects removed. The vertical dashed line marks the day of the tax change, July 1, 2015.

Figure 4: Processor After-Tax Prices and Revenue Bandwidth Sensitivity

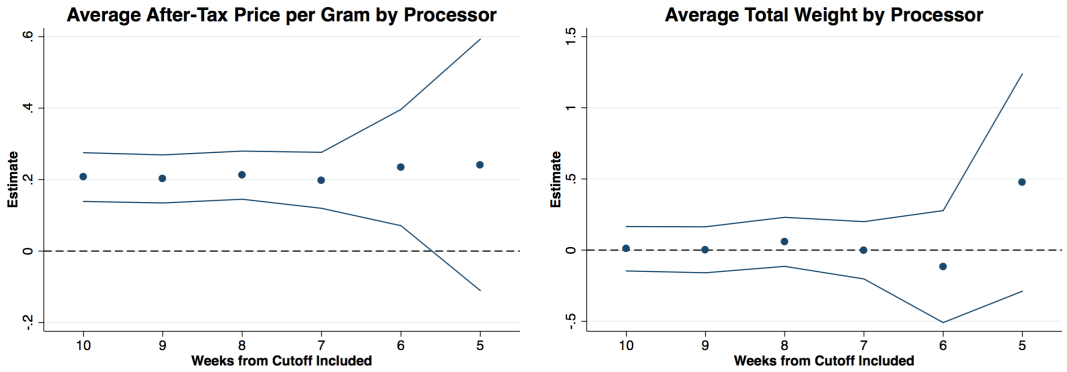

These figures consider the sensitivity of the estimates in Column (1) of Table 8 to the number of weeks of data we include on either side of the tax change. The dots mark the estimates for each bandwidth choice and the lines mark the 95% confidence intervals around these estimates.

Figure 5: Vertical Integration

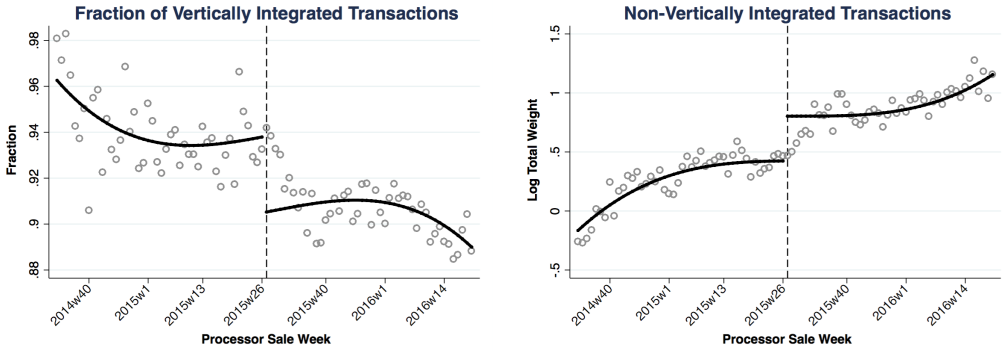

These figures are based on the estimates in Column (1) of Table 10. The solid line plots the estimated response to the tax change plus the polynomial. The scatterplot is the corresponding weekly average of the dependent variable with estimates of the day-of-week, day-of-month, and processor fixed effects removed. The vertical dashed line marks the day of the tax change, July 1, 2015.

Figure 6: Vertical Integration Bandwidth Sensitivity

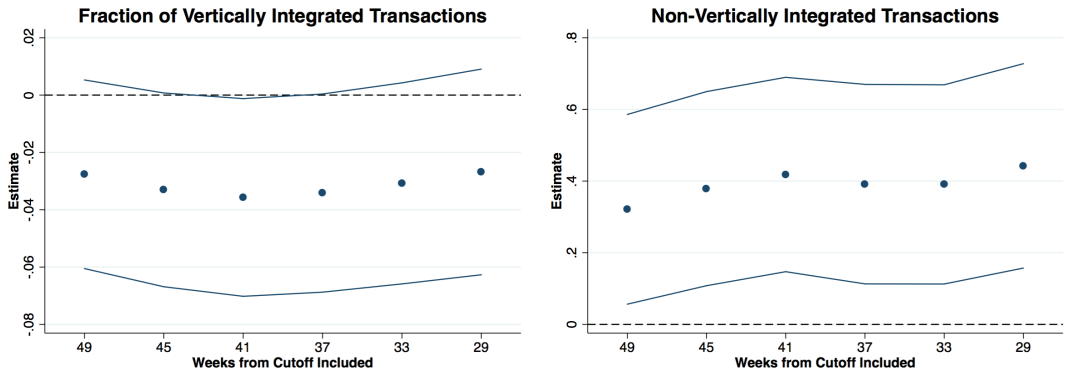

These figures consider the sensitivity of the estimates in Column (1) of Table 10 to the number of weeks of data we include on either side of the tax change. The dots mark the estimates for each bandwidth choice and the lines mark the 95% confidence intervals around these estimates.

Figure 7: Retail Tax-Inclusive Prices and Weight

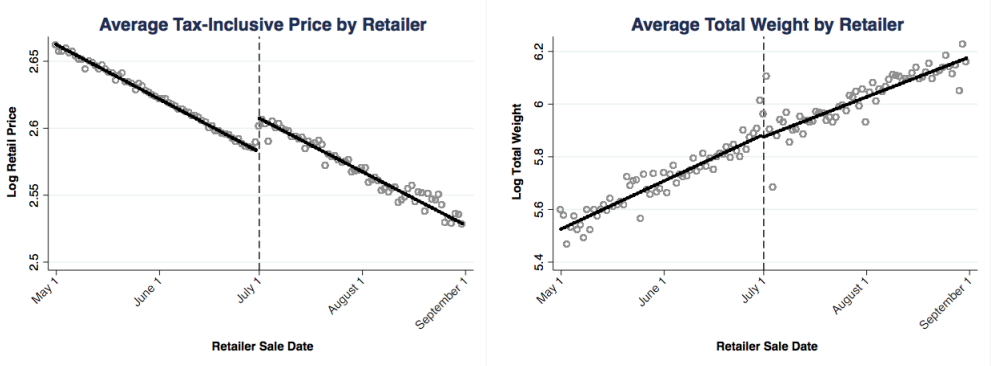

These figures are based on the estimates in Column (1) of Table 13. The solid line plots the estimated response to the tax change plus the date polynomial. The scatterplot is the corresponding daily average of the dependent variable with estimates of the day-of-week and day-of-month fixed effects removed. For the figure on the left, inventory-lot fixed effects were also removed. For the figure on the right, retail fixed effects, log processor price, log competitors' processor price, and whether any competitors were also removed. The vertical dashed line marks the day of the tax change, July 1, 2015.

Figure 8: Retail Tax-Inclusive Prices and Weight Bandwidth Sensitivity

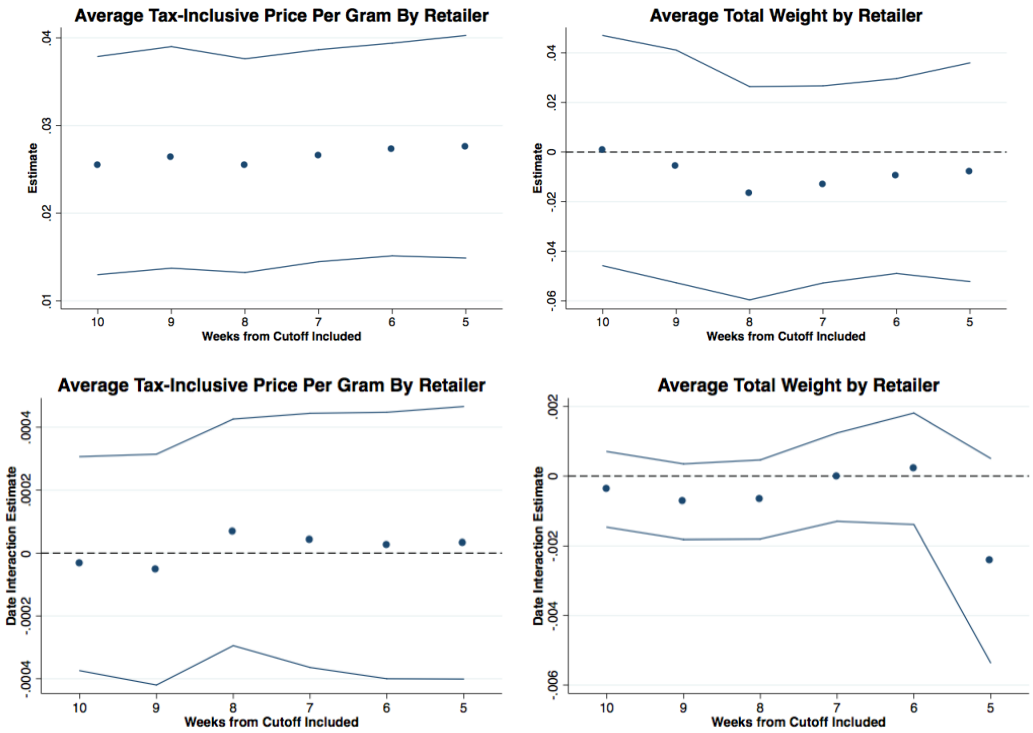

These figures consider the sensitivity of the estimates in Column (1) of Table 13 to the number of weeks of data we include on either side of the tax change. The dots mark the estimates for each bandwidth choice and the lines mark the 95% confidence intervals around these estimates.

Figure 9: Retail Tax-Inclusive Prices and Weight Polynomial Choice Sensitivity

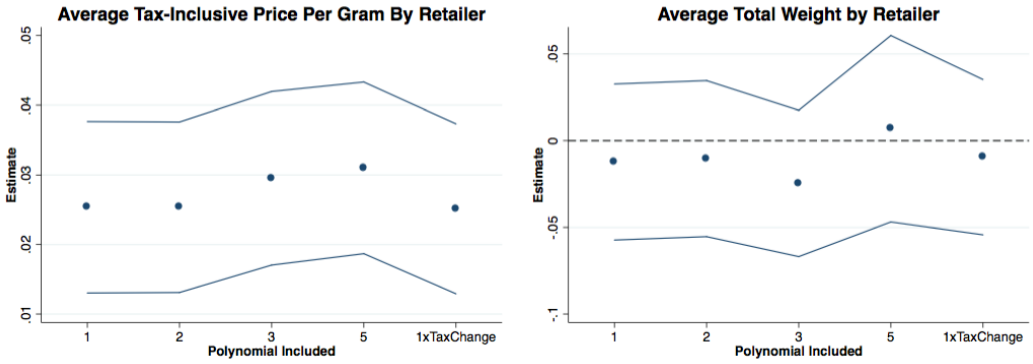

These figures consider the sensitivity of the estimates in Column (1) of Table 13 to the polynomial order chosen. We consider linear, quadratic, cubic, quintic, and linear interacted with the tax change indicator. The dots mark the estimates for each polynomial choice and the lines mark the 95% confidence intervals around these estimates.

Figure 10: The Average Price of One Gram of Marijuana and Tax Incidence across Markets

In this figure, we plot the average retail firm's price of one gram of marijuana both before and after the tax change. We then consider how much goes to processor and retail taxes as well as how much is spent to purchase a gram, on average, from the firm. Before the tax change, all prices and taxes (in dollars) are based on the average prices the month prior to the tax change. After the tax change, these numbers are the pre-tax change prices adjusted by our estimated changes caused by the tax changes. This holds constant the composition of the market and eliminates any secular trends in prices.

Appendices

A Incidence with the Federal Income Tax

In this appendix, we consider how our incidence results would be affected if we incorporated the federal personal income tax into our calculations. Without considering the federal income tax, we found that consumers bore about 44% of the increased taxes due to the tax reform on July 1, 2015. Understanding exactly how federal income taxes changed in expectation in response to the tax reform is challenging. The wording of the original Washington law left open the possibility that retail firms could owe federal income taxes on the tax-inclusive price.[33] This was part of what instigated the reform of the tax system on July 1, 2015. An IRS memo (Moffitt 2015) was eventually written, which clarified that federal taxes were not owed on taxes paid to the state (effective retroactively), but this memo was not issued until July 31, 2015, after the tax reform had taken place.

If retailers anticipated this decision by the IRS, then there was no change in the amount of federal income taxes owed other than that caused by the change in retail and processor prices. If retailers did not anticipate the decision, the tax reform would have appeared to provide a large federal tax break to retailers—before the tax change, they owed federal taxes on the tax-inclusive price, and after the tax change, they owed federal taxes on the tax-exclusive price. While we cannot say exactly what incidence looks like when we incorporate federal personal income taxes without knowing all firms' expectations, we examine two benchmark cases. First, we consider what incidence would look like if firms always believed that federal taxes would apply to the tax-exclusive price. We then consider what beliefs would have needed to look like in order for taxes to be completely passed-through to consumers as an alternative benchmark.[34]

We do not directly observe federal income taxes owed in our administrative data, so we must estimate them. Section 280E—the federal tax code that governs the taxation of illegal substances—and subsequent Chief Counsel Advice (CCA 201504011) specifies that retailers are only allowed to deduct the costs of goods sold (as opposed to, say, labor and capital

[33]The original law stated: "This tax is the obligation of the licensed marijuana retailer, is separate and in addition to general state and local sales and use taxes that apply to retail sales of tangible personal property, and is part of the total retail price to which general state and local sales and use taxes apply." The last part of this sentence led some retail firms to believe they owed federal income taxes on their tax-inclusive price.

[34]We do not consider processor federal income tax liability in this section for two reasons: (1) it is much less likely that processors thought they were subject to a federal income tax on their tax-inclusive price before the tax change, so there would be no real change in the amount of federal income tax owed as a result of the reform except due to changes in the processor price, and (2) cultivators and processors are allowed more deductions under federal law, many of which we do not see, so our calculations of the federal income tax liability would be much less accurate.

costs), which in this case consists of the price paid to the processor and any transportation costs of obtaining the inventory. We thus calculate a retail firm's federal income tax liability assuming the following: (1) they are not organized as a C-corporation (so they will file the personal income tax on their business income, not the corporate income tax), (2) their only income source is their marijuana business, (3) their only cost of goods sold is the marijuana purchased, (4) they do not itemize, and (5) they are married with one child. We recognize these assumptions are not perfect, but they are the best we can do given what we observe and reasonable variations in any one of these would not substantially change the estimated tax liability on a gram of marijuana.

When all firms believe federal income taxes are always assessed on the tax-exclusive price, we calculate how federal income taxes changed on a dollar of marijuana at the new equilibrium price, \$13.49, but we allow the processor price to change with the tax reform. We find that at the new processor price, federal income taxes on a dollar of marijuana are \$1.54 and at the old processor price are \$1.48; that is, federal income taxes on a dollar of marijuana rose by 6 cents, which amounts to a 0.44% increase from the equilibrium tax-inclusive price, \$13.49.[35]. Hence, the net change in taxes and marginal costs is 5.50%, so consumers bear 42% of the tax burden (=.0233/.0550).

As an alternative benchmark, we consider what happens when 70% of firms believe that federal income taxes are always assessed on the tax-exclusive price and 30% of firms believe that federal income taxes are assessed on tax-inclusive prices before the tax reform and tax-exclusive prices after the reform. We calculate how federal income taxes changed on a dollar of marijuana at the new equilibrium price, \$13.49, but we allow taxation to switch between tax-inclusive and tax-exclusive prices and we allow the processor price to change with the tax reform. In this scenario, federal income taxes are \$1.90 (=.3*\$2.74+0.7*\$1.54) on average before the tax reform. Federal income taxes on a dollar of marijuana declined by 36 cents, which amounts to a 2.7% increase from the equilibrium tax-inclusive price, \$13.49.[36]. Hence, the net change in taxes and marginal costs is 2.33%, so consumers bear 100% of the tax burden (=.0233/.0233). As the fraction of firms that expected to pay taxes on the tax-inclusive price before the tax change increases above 30%, pass-through to consumers would be more than 100%. And, as it decreases below 30%, pass-through to consumers would be less than 100%.

[35]We calculate this as the log change, $0.0044 = log((13.49 + .06)/13.49)$.
[36]We calculate this as the log change, $0.027 = log((13.49 - 0.36)/13.49)$.