

NBER WORKING PAPER SERIES

GETTING INTO THE WEEDS OF TAX INVARIANCE

Benjamin Hansen
Keaton Miller
Caroline Weber

Working Paper 23632
<http://www.nber.org/papers/w23632>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2017, Revised January 2021

The authors would like to thank David Agrawal, Nathan Anderson, Youssef Benzarti, David Evans, Michael Grossman, Bill Hoyt, Donald Kenkel, Michael Kuhn, Nathan Seegert, Joel Slemrod, and Dave Wildasin for helpful comments. We appreciate comments and feedback from participants at seminars at Case Western, Columbia, Cornell, University of Kentucky, Portland State, Rutgers, and conference participants at IHEA, NTA, WEAI, and the IIOC meetings, as well as industry participants and Cannabis Science and Policy Summit attendees. Many thanks to David Shi for excellent research assistance. This paper previously circulated as part of “The Taxation of Recreational Marijuana: Evidence from Washington State” and some of our thanks are for comments provided on that earlier draft. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Benjamin Hansen, Keaton Miller, and Caroline Weber. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Getting into the Weeds of Tax Invariance
Benjamin Hansen, Keaton Miller, and Caroline Weber
NBER Working Paper No. 23632
July 2017, Revised January 2021
JEL No. H2,H20,H21,H22,H23,H25,H26,H32,H71,I1,I18,K4

ABSTRACT

We provide the first general empirical test of tax invariance (TIV). When a 25 per- cent tax remitted by manufacturers was eliminated in Washington state and the retail cannabis excise tax was simultaneously increased from 25 to 37 percent—a shift intended to be revenue-neutral—TIV did not hold. Manufacturers kept two-thirds of their tax savings instead of passing all their savings through to retail firms via lower prices as predicted by TIV. One-third of the retail tax increase was passed on to consumers via higher retail prices—TIV would have predicted constant or even declining tax-inclusive retail prices.

Benjamin Hansen
Department of Economics
1285 University of Oregon
Eugene, OR 97403
and NBER
bchansen@uoregon.edu

Caroline Weber
University of Kentucky
423 Patterson Office Tower
Lexington, KY 40506
caroline.weber@uky.edu

Keaton Miller
University of Oregon
Department of Economics
1285 University of Oregon
Eugene, OR 97403-1285
keatonm@uoregon.edu

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

Taxation and revenue generation has been one of the primary political and economic arguments for legalizing recreational marijuana (Miron and Zwiebel, 1995). However, little is known about the behavior of participants in legal marijuana markets or, in particular, how industry participants may respond to alternative policies - particularly given the existence of readily available substitutes in the form of sophisticated black market goods and legal or quasi-legal medical marijuana. Policy makers across states have implemented widely different regulatory regimes: Colorado requires retail firms to grow 80 percent of the product they sell while Washington forbids retailers from growing at all. Washington's current tax rate is 37 percent; Massachusetts has a tax rate of 3.75 percent.² In addition to the difference in tax rates, states apply taxes at different points in the supply chain. These differences point to substantial variation in beliefs among policy makers 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 for purchase at retail stores on July 8, 2014. Washington's regulations created a supply chain with three

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

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

types of firms: cultivators, which grow the cannabis plant, processors, which transform the raw plant material into usable marijuana³, and retailers, who sell products to end-users. To comply with state law, firms must provide detailed information on their operations to the state government. These unique “seed-to-sale” administrative records track the entire supply chain, and are verified by the government through frequent audits. We obtain this data, which gives us the unique ability to observe variation in marijuana prices, quality and variety as firms enter and compete in the marketplace. Previous studies of the marijuana industry have examined the illegal market through surveys, prices collected from drug seizures, or crowd-sourced data (Jacobi and Sovinsky, 2016; Miron and Zwiebel, 1995; Pacula et al., 2000; Clements and Zhao, 2009; Donohue, Ewing, and Peloquin, 2011; Williams, van Ours, and Grossman, 2011), and some have used this evidence to speculate about legal market outcomes. We are the first to examine the legal market directly.

To understand the effect of market regulations and taxation on industry outcomes, we take advantage of a reform to Washingtons marijuana tax policy. Prior to July 1, 2015, a 25 percent gross receipts tax was assessed at each transfer of marijuana (i.e. between cultivators and processors, processors and retailers, and retailers and consumers).⁴ After July 1, 2015, the only tax collected was a 37 percent excise tax at retail. Crucially, this change was unexpected by market participants: the reform legislation was passed during a special session of the Washington Legislature on June 27, 2015, and signed by the Governor

³Here and throughout, we consider only ‘usable marijuana’, which is defined as dried and cured flowers for consumption by smoking. Different ‘products’ refer to different strains of marijuana, which are analogous to the different cultivars of apples commonly available in supermarkets. While other derivatives of the cannabis plant, such as oils, topical creams, and edible products, are available, usable marijuana comprises the vast majority of the market by transaction count, and revenue.

⁴This was implemented as an excise tax at each transfer of marijuana.

on June 30.⁵ We estimate the effect of these policy changes on key behaviors and decisions made throughout the production process, while controlling for product quality and other factors that may impact demand and supply decisions.

We find the original tax regime encouraged vertical integration between cultivators and processors (Washington bans vertical integration between retailers and other firm types). Due to the variety of business relationships observed in the data, we define vertically integrated marijuana as product for which the cultivator and processor were the same firm.⁶ We capture changes in vertical integration by examining two measures: 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 percent of wholesale transactions comprised vertically integrated marijuana. This drops to about 90 percent following the tax change. The weight of dis-integrated marijuana sold to retailers increased by 42 percent after the change.

Taken together, these results suggest the original gross receipts tax created inefficiencies in the supply chain that persisted after the removal of the taxes due to fixed costs. 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 the drop in the percentage of vertically integrated transactions quadruples to 16 percent, and the increase in the weight of dis-integrated marijuana more than doubles to 105 percent.

The tax-invariance folk theorem makes the claim that if taxes remain the same, but a tax

⁵Contemporaneous 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).

⁶Vertically integrated processors may process marijuana produced by other cultivators.

reform changes where in the supply chain they are assessed, after-tax prices at each point in the supply chain will remain the same. Since we find 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 policy-makers tasked with setting taxes, the analysis is complicated by the degree of product differentiation – the typical retailer has many tens of 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. Taken together, these results suggest 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 the tax-inclusive price faced by consumers for identical products increased by 2.3 percent. 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 percent of the retail tax burden. We find the quantity purchased decreased by 0.95 percent, 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 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: higher taxes on the margin may not increase revenue.

Our vertical integration findings contribute to several broader long-standing public finance questions. It is theoretically obvious that vertical integration would be a natural consequence of a gross receipts tax and tax economists frequently come out vehemently against gross receipts taxes for this and other reasons (e.g., McClure, 2005; Pogue, 2007; Testa and Mattoon, 2007); however, this is the first paper we are aware of that provides compelling empirical evidence of this behavior. Gross receipts taxes have begun to proliferate across states in recent years, so this paper 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-collection invariance along two dimensions (Kopczuk et al. 2016; Chetty, Looney and Kroft, 2009). First, we highlight that the possibility of vertical integration under a gross receipts tax is another reason tax-collection invariance will no longer be expected to hold. Second, we document that tax collection at the retail level has different incidence implications than does tax collection at the processor level. In contrast to previous work, there is no evidence in this case that the result is driven by different tax evasion possibilities at different points in the supply chain (Kopczuk et al. 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 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, among others, which generally finds consumers bear the substantial majority of the tax burden (Kopczuk et al., 2016; Harding, Liebtag, and Lovenheim, 2012) and can exceed 1 (Kenkel, 2005; Barnett, Keeler, and Hu, 1996). 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 could also be explained by the decline in federal income taxes some retail firms expected to pay upon the switch from a gross receipts tax to an excise tax.⁷

Finally, our findings additionally contribute to broader research on supply side interventions in drug markets, which has included examinations of policies ranging from alcohol’s prohibition (see Miron and Zwiebel, 1991), to limitations on precursors to methamphetamines (see Dobkin and Nicosia, 2009) and the legalization of marijuana sale and cultivation for medical purposes (Anderson, Hansen and Rees, 2013). Our research is some of the first to study legalized marijuana markets for recreational use, and the first to analyze the substantial natural experiment the tax change offers. And while prior papers have focused on the elasticity of demand for illegal markets, we are the first, to our knowledge to study the market-level elasticity of demand in a legal marijuana market.

The remainder of the paper proceeds as follows. In Section 2, we discuss the history

⁷The 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 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. We provide more details on this in Section 4.3.

of marijuana’s legalization and the associated 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 by discussing the policy and economic implications of our findings, such as where Washington’s tax rate is on the Laffer curve and vertical integration’s role in the efficiency of these newly legalized markets.

2 Background

Prior to 1938, marijuana was a legal substance 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 via the Controlled Substances Act of 1970 significantly strengthened the prohibition against marijuana, as it was quickly classified as a Schedule I substance with a ‘high potential for abuse and little known medical benefit.’⁸

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 (including Washington D.C and Puerto Rico) permit the cultivation and use of marijuana for medical reasons. In response to the growing acceptance of the medical use of marijuana, in October 2009, the United States Department of Justice issued a memorandum to United States Attorneys (Ogden, 2009) discussing the appropriate way to

⁸Other Schedule I substances include heroin and methamphetamine.

allocate resources in states with legal medical marijuana markets. In particular, the memo stated “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.” This was broadly interpreted as an effort to defer to states’ choices in the absence of federal consensus (Stout and Moore, 2009), though the Department of Justice emphasized a need to investigate and prosecute “drug traffickers who hide behind claims of compliance with state law.”

In the November election of 2012, Washington voters approved Initiative 502 which legalized the recreational possession and consumption of small amounts of marijuana for adults over 21. Three types of licenses were created: producers (who we term ‘cultivators’ in this paper), who are permitted to grow and harvest cannabis plants, processors, who transform the harvested plant material into usable marijuana and other products for wholesale, and retailers, who may sell final products obtained at wholesale to consumers. The Initiative specified that taxes would be collected with the revenue set aside for education, healthcare, and substance abuse prevention. The initial tax structure assessed a 25 percent tax on each marijuana transaction. This includes when grown marijuana is sold to processors who convert the harvested plant material into usable marijuana⁹, when processors sell the usable marijuana to retailers and when retailers sell it to end consumers.

⁹‘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. Though the traceability system has a unique code for products that fall into this category, it contains two types of products: both raw dried flowers and pre-rolled joints, which include some value-add. Pre-rolled joint products are also listed under a different inventory type code. As we cannot cleanly distinguish between raw flower and pre-rolled joints in the “usable marijuana” category, our analysis includes the entire “usable marijuana” category as well as any pre-rolled joint products we can identify across other categories. 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.

Cultivator licenses come with capacity constraints – each cultivator is allocated one of three sizes of plant canopy¹⁰ Firms which grow cannabis plants may also have a license to process the material into usable marijuana, and vice versa. These vertically integrated firms did not owe any taxes when they transferred the grown marijuana to their processing operation. However firms involved in the production or processing of marijuana are forbidden from owning and operating a retail location. Holders of a retail license were initially limited to three locations. Broader ownership restrictions exist as well – individuals are not allowed to have direct or familial financial interests in more than one marijuana license.

The initiative gave regulatory authority to the newly-renamed Washington State Liquor and Cannabis Board. One of the early actions the new board took was the implementation of a state-operated traceability system which would track the cultivation, testing, processing, and retail sales of marijuana throughout the state. We provide more details on these administrative data in Section 3.

The federal Department of Justice responded to the changing environment in Washington and Colorado¹¹ with an August, 2013 memo – commonly known within the industry as 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

¹⁰According 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.”

¹¹Colorado voters also approved a legalization effort in November, 2012. We focus on Washington in this paper for two reasons. First, in contrast to Washington’s vertically separated market, Colorado’s system enforces mandatory vertical integration – each retailer sells product they have grown individually. Secondly, and perhaps more critically, while both states collect similar traceability data, Washington’s legal framework requires public disclosure of the data and Colorado’s forbids it.

about 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 set a clear expectation that “states and local governments ... will implement strong and effective regulatory and enforcement mechanisms.” Within the industry, the traceability system is seen as the implementation of the expectations and priorities laid out within the Cole Memo.

The tax reform analyzed in this paper was part of House Bill 2136 introduced during the 2015 Regular Session of the Washington Legislature. As our identification rests on the assumption that the policy change was effectively a set of exogenous cost shocks throughout the supply chain, the details of the bill’s history are critical. Table 2 provides a detailed 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. However, while a Senate committee recommended passage on the last day of the session, 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. It wasn’t until the very end of the Second Special Session, June 27, that the bill received a vote by the full Senate. 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, which 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. Notably, Washington 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 almost certainly 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.

3 Data and Methods

Our data consist of administrative records from the Washington State Liquor and Cannabis Board (WSLCB). As a part of the legislative effort to legalize and regulate recreational marijuana production and consumption in Washington, the state implemented a traceability system (also known as a seed-to-sale system) produced by BioTrackTHC. Its purpose is to track each step in the marijuana supply chain, enabling regulators to collect taxes and prevent diversion to the black market. Compliance is enforced through random inventory audits, backed by penalties ranging from fines to inventory seizure and destruction. The state provides an API to market participants and requires timely reporting of many details of the production process. Firms generally use one of several commercially available software packages to report data to the traceability system and comply with state data-reporting requirements. The end result is data that tracks a cannabis plant’s planting, harvest, production 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 foreign contaminants and moisture content to ensure quality and safety, as well as for potency via the concentration of tetrahydrocannabinol (THC), tetrahydrocannabinolic acid (THC-A), and cannabidiol (CBD), the primary psychoactive components of the cannabis plant.

We utilize an extract of the state’s database that includes all plants and products but removes information about the supply chain that is subject to security and privacy concerns.¹² 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

¹²The raw data come in the form of SQL tables which do not include data on individual employees. Additionally, while firms are required to report itineraries and planned routes for marijuana transfer operations (e.g. when a wholesaler makes a delivery to a retailer), our extract does not include them.

a seed) as well as the strain¹³ and a log tracks its movement through the growth process. Once harvested, flowers and other plant material are generally collected and converted into a new inventory lot which is assigned a unique identifier to represent the homogeneity of product within an inventory lot. These intermediate products may progress through several processing steps including conversions, combinations, and divisions before wholesale. Along the way, the database records these transformations, generating new inventory lot ids as necessary.

The last step of production is the division of a large wholesale inventory lot into multiple smaller lots for sale to individual retail stores. Each lot consists of multiple sealed packages of a specific weight of marijuana (most commonly 1 gram, 2 grams, or 3.5 grams) which are considered identical. When lots are sold to retail locations, the tracking system records the date, quantity and price of the transaction and assigns a new id. As a consequence, the observation of an inventory lot id at the retail store uniquely identifies the retailer, processor, and cultivator, as well as the strain 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,¹⁴ in a retail transaction to the relevant inventory lots, allowing us to trace sales back through to a set of original plants.¹⁵

¹³Strains are defined by the cultivator and consequently are the “dirtiest” field in our data. Additionally, the system allows cultivators to input the seed-grown offspring of two cloned parents as a member of that clone strain, even though the offspring may not ‘breed true.’

¹⁴Washington’s regulatory framework 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. While the order of transactions is maintained and other information in the system (specifically the inventory log file) may allow us to track the specific time of each transaction, we choose the daily level as our most granular view of activity in the industry.

¹⁵Given the details of the production process discussed in this section, it is not possible to precisely identify which plant or plants a particular retail package of marijuana came from – we can only identify the set of plants which contributed to the creation of the wholesale lot from which the retail package was

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.¹⁶ Given the retail firms’ inability to access traditional financial services markets, 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 a consequence, the posted prices faced by consumers include all taxes and therefore, in contrast to many other settings, the sales taxes are salient to consumer-level decision making.

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. For each component of the supply chain—cultivators, processors, and retailers—we collapse the data by firm-day unless otherwise specified after performing minor cleaning steps, none of which significantly change our results.¹⁷ We provide additional details about data cleaning in the Online Data Cleaning Appendix.¹⁸

Our primary target in our regression analyses is the response to the tax change. As the change took place within the broader context of the market’s non-linear evolution, our esti-

divided.

¹⁶The 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-exclusive prices. To clean the price data, for each firm, we search for a multiplier within the bound created by the maximum effective tax rate to maximize the percentage of prices with round numbers. We verified the results of this algorithm using public-facing menu data captured by the Internet Archive. Additional details on our price-cleaning process are given in the Online Data Cleaning Appendix: <http://www.keatonmiller.org/s/data-cleaning-appendix.pdf>

¹⁷We drop a small number of duplicate sales in the processor data and we drop all retail transactions that were deleted or refunded. We also drop all processor and associated retail transactions in which the processor price or total transaction amount was less than or equal to one cent. It is our understanding that these were generally samples given to the retailer by the processor; they were given a price of one cent so that these movements of marijuana could be tracked within the traceability system.

¹⁸Data cleaning appendix available at <http://www.keatonmiller.org/s/data-cleaning-appendix.pdf>

mating equations include a polynomial in time. Furthermore, as the change was unexpected by market participants, we include several dummy variables in 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. The detail of the data collected by the traceability system allows us to analyze behaviors at the firm-day level, though cyclical patterns throughout the course of a week and a month cause 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:

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

where y_{it} is our outcome variable for firm (cultivator, processor, or retailer) i at date t , tax_{it} is an indicator variable that is one after the tax change took place on 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, and $date_{it}^l$ is a polynomial in the date of sale. June 30 is the date the law was signed, and July 4th is a holiday which resulted in deviations from normal trends. We explore the sensitivity of our estimates by shrinking our adjustment window to only include June 30 and July 1, or extending the adjustment window to June 27th (the date the tax reform was passed by the Washington Legislature). We also explore the sensitivity of our analysis to the inclusion of day-of-month fixed effects based on the hypothesis that these

may not always be very important, but do consume many degrees of freedom. Lastly, we consider the sensitivity of our estimates to changes in the polynomial we use as well as the inclusion of county-specific polynomials to allow for county-specific time trends.

We take the logs of all of our outcomes, unless otherwise specified, because the outcomes we examine 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 the details of our estimating equations, we interpret coefficients on the tax change indicator variable to represent the average short-term response to the tax reform unless otherwise specified.

4 Results

We report findings as marijuana flows through the supply chain: we first examine cultivators, then discuss processors, and finish with the retail market. Across markets, we present tables consisting of point estimates for each of the outcomes we examine, and table and figures of robustness checks for outcomes of particular significance. For these same important outcomes, we also provide figures detailing the data and graphically plotting the response. In the figures, the solid line is a fitted model based on 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 estimate of the day-of-week, day-of-month, and firm location fixed effects. We also remove the estimates on the indicators for June 30 - July 5 from the fitted line.

4.1 Cultivators

We first focus on cultivators – those firms that plant and harvest raw flowers from the cannabis plant. Table 3 provides summary statistics for cultivators. On average, firms harvest plants 116 days after planting. Vertical integration is rampant, as over 95 percent of growers also have a license to process marijuana. The July 1 policy reform decreased their transaction tax from 25 percent to zero. However, because of the rampant 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.¹⁹ We thus focus on other outcomes, which, while they may not be responding to a meaningful change in the cultivator tax because of vertical integration, may still respond to a decrease in their vertically integrated processor’s transaction tax from 25 percent to zero. 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 no plantings or harvests took place that day at that firm. The left panel of Figure 1 illustrates the the counts of the plants harvested around the July 1 tax change. The harvest rate increases by 7.5 percent, but this change is not statistically significant. However, we may be underpowered

¹⁹The administrative data does not contain meaningful prices for many of the transfers of plant material between cultivators and processors. In the administrative data, many of the processor prices are missing or are filled in with prices arbitrarily close to zero. Due to the difficulty of interpreting these values, we focus on quantity outcomes.

to detect effects in the harvesting and planting data because these activities are infrequent and when they happen occasionally happen in large volumes. We examine the the number of days from planting until harvesting, and find evidence that the number of days from planting until harvest falls by 6.9 percent. 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. The empirical results are largely insensitive to bandwidth, allowing for an adjustment period from June 27 through July 5, 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. In some ways, given the vast amount of vertical integration and that firms are producing at capacity due to strong demand, this is not entirely surprising – particularly given the firms did not have a long period to adjust their production process since the law was only passed days before the tax change went into effect.

4.2 Processors

The tax reform also eliminated the transaction taxes faced by processors – those firms which take raw marijuana flowers and plant material as inputs and transform them into usable marijuana and other products. Before the change, the processors paid the taxes after the transaction, implying the retailers paid the equilibrium 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.²⁰ The only significant findings are on prices – we find that average equilibrium prices fell by 8 percent. Simultaneously, the after-tax price (the price received by processors after taxes) increased by 21 percent.²¹ We also find the quantity of sales transactions, total weight of marijuana sold, and THC levels were unchanged.²²

Figure 3 provides graphical evidence of the after-tax price and weight responses we estimated in columns (2) and (3) of Table 7. Because equilibrium prices did fall, but the after-tax price for processors increases dramatically following the tax change, it appears that the elimination of the 25 percent tax was a huge boon to processors, while it simultaneously reduced the portion of variable costs of retail firms attributable to the cost of goods sold. At the same time, we find the quantity of marijuana processed and sold to retail firms did not significantly change around the reform. This suggests two things. First, it provides evidence the 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 the quantity of marijuana substantially decreases (indeed June 30 appears to be a true outlier, consistent

²⁰The 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.

²¹The log change in the tax rate is 29% (i.e. $\log((1-.75)/1)=0.287$).

²²Within 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 the “usable marijuana” category and thus outside our analysis). Our null result here suggests firms did not substitute in this way.

with tax avoidance on the part of 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, polynomial order, interacting the polynomial with county indicators, and exclusion of day-of-month indicators.

In summary, our findings suggest processors widely benefited from the elimination of their transaction taxes with retailers. The tax benefits are shared between both the retailers, whose price paid for marijuana fell by 7 percent on average, and the processors, whose price received after paying taxes rose by 21 percent. After we estimate the retail response to the tax regime change in Section 4.3, we will discuss the tax invariance and incidence implications of these results. We find no significant increase the quantity of marijuana supplied from processors to retailers, other than a one day anticipation effect and a few days adjustment period (in which the sales withheld the day before the tax are sold to retail firms after the tax change). This suggests that in the short run, the tax changes did not fundamentally alter the supply of marijuana. This altogether suggests that in the short-run window surrounding the tax change, supply is relatively inelastic – not surprising given the requirements of the production process.

4.2.1 Vertical Integration of Cultivators and Processors

In this section, we provide the first empirical evidence on how vertical integration behavior adjusts to an 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 (e.g. marijuana cultivators are not also selling to cigarette companies, which we do not observe).

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 (even if those firms were formally vertically integrated).

Before the tax reform, cultivators and processors had a strong incentive to vertically integrate because there was a 25 percent gross receipts tax on cultivators that could be avoided if the cultivator sold their product to a vertically integrated processor. Anecdotally, cultivators often wished they did not also have to invest in the equipment to be processors, but could not generally afford to do otherwise. Moreover, even if a company was both a cultivator and processor, they would have still liked to sell some of their raw product to other processors or purchased some raw product from other cultivators, but both were often made too expensive by the existence of the 25 percent tax. These stories are borne out in the data as about 93 percent of all inventory lots sold by processors to retailers consisted of vertically integrated marijuana. Additionally, only about 5 percent 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 on July 1, 2015 eliminated the incentive to vertically integrate for tax reasons. We don’t expect many firms would de-integrate and become only a cultivator or a processor

because of the fixed costs already paid, 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 understand the change in behavior resulting from the tax reform, 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 roughly six weeks after processors purchase raw material from cultivators before the resulting products are sold to retailers. This leads us to make two adjustments to our analysis in this section relative to our previous processor analysis: (1) we aggregate the analysis at the weekly level, and (2) we include indicator variables for the six weeks following the tax change to examine the adjustment period, relative to the long-run effect. In contrast to the rest of the processor analysis, we also do not drop new firms because these are of particular interest as they may be the 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 transactions which did not originate from a vertically integrated cultivator. Point estimates are provided in Table 9. The fraction of vertically integrated sales falls by 3.7 percent after the adjustment period (Column 1), which is driven by a 42 percent 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. Table 9 Column (3) 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 con-

sidering the behavior of each firm in its first week in business – i.e. 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 percent and the weight of vertically dis-integrated marijuana increases by 105 percent after the reform. Table 10 and Figure 6 illustrates the robustness of the results controlling for trends in a variety of ways, adjusting for shorter or longer adjustment periods, and variation in bandwidth.

In summary, we find in the long-run, that the volume of non-vertically integrated transactions rose substantially, although vertically integrated transactions continue to dominate the market in absolute size. Most of the estimates in this section are likely lower bounds on the true inefficiencies of a gross receipts tax because we are examining the elimination of a gross receipts tax, and so many firms will have already paid the fixed costs necessary to become both cultivators and processors, and may not choose to become one or the other afterwards, even if that would have originally been efficient in the absence of a gross receipts tax.

4.3 Retail Market

We now examine the consumer retail market. The tax reform converted the tax at the point of retail sale to an excise tax which, given the way the original law was written, switched the sales tax from applying to the tax-inclusive price before the tax change to the tax-exclusive price after the tax change. The reform also increased the tax rate from 25 percent to 37 percent. Given that we can observe a portion of any given retailer’s input prices and those

²³The responses we observe for all firms is not driven by firms in their first week in business.

of their nearest competitors, we add some new covariates to our regression analysis relative to previous sections. In particular, we add the prices paid to processors (their variable costs), the prices of local²⁴ competitors, and an indicator for whether there were any local competitors. This has the potential to be valuable for the price estimates because, at the firm level, the estimated price could change not only due to a change in the prices, but also due to a change in the composition of goods sold. The inclusion of these variables is intended to control for these changes. Retailers also have significant weekend patterns in weight sold and these can vary substantially by firm depending on their location within a city as well as the days of the week they are open. Hence, we include interactions between each firm indicator and the indicators for Friday, Saturday, and Sunday in all our baseline estimates; this increases the precision of our estimates. 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. We find evidence that tax-exclusive prices, the prices retailers received, fell by 4.8 percent. We find tax-inclusive prices, the prices paid by consumers, rise by 2.1 percent.²⁵ We find no significant response in weight or sales to this marginally significant increase in the prices consumers paid. We find a statistically significant decline in THC content which, when combined with the null results for THC in the processor analysis, suggests consumers engaged in a degree of product-level substitution

²⁴We define ‘local’ as ‘within 10 miles.’

²⁵The excise tax increased from 25 to 37 percent 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 percent 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 percent.

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 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 an inventory lot is fixed. We also anticipate our estimates will be substantially more efficient when we estimate the response at the inventory-lot level because we are able to more effectively control for quality via inventory-lot fixed effects. We do not estimate the weight response at the inventory lot level because we do not have the same substitution problem in our firm-level estimates. If individuals shift away from one product towards another, this will leave our firm-level weight estimates unchanged. It could, however, become more problematic if we estimated the response by inventory lot because we are estimating the log change and an increase of one unit and a decrease of one unit are not equivalent unless the sales volume of the two inventory lots started 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 percent and this estimate is significant at the one percent level. The magnitude of the inventory-lot and firm level estimates are quite similar, suggesting that

²⁶Our analysis at the inventory-lot level is restricted to all inventory lots that sell at least once in the month prior to the tax change. If we did not impose this restriction, our estimates would still be identified off of inventory lots that sell both before and after the tax change; however, including this restriction allows us to more accurately estimate the coefficients for the inventory lots from which our identification comes.

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

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 tax-inclusive price (estimated at the inventory lot-day) and weight (estimated at the retail sales location-day) 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 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 and it allows us to conserve degrees of freedom and assess whether the tax change also affected the steady decline in prices or increase in weight in this market over time.²⁸ This short-run price elasticity of demand suggests Washington is on the part of the Laffer curve where higher taxes on the margin would increase revenue, and this is consistent with the positive effect on tax-inclusive revenue in Table 12 Column (6). However, the decrease in *tax-exclusive* prices 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 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 don't have the power to detect this in our estimation, our visual inspection of the raw data in Figure 7 suggests that the decline in the growth rate of marijuana demand seems to slow down several weeks after the tax change, suggesting that this is about when we hit the medium-run elasticity of demand after the tax change, holding all else constant. This larger medium-run elasticity of demand suggests that Washington is

²⁸See Figure 9 for evidence that the estimates are not sensitive to polynomial choice.

near the peak of the Laffer Curve in the medium-run.

We consider sensitivity of our main price and weight estimates in Table 13 Columns (2) to (7) and Figures 8 and 9. Our estimates are insensitive to bandwidth, allowing for an adjustment period from June 27 through July 5, decreasing the adjustment period to June 30 - July 1, interacting the polynomial with county indicators. We also show 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 dollar amount spent to purchase a gram of marijuana from the retailer, 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 numbers. 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 was \$3.46 (\$1.03 collected and paid by the processor and \$2.43 paid by the retailer) and the average taxes paid on a gram of marijuana after the tax change was \$3.42 (all collected and paid by the retailer).²⁹ 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. However, we find that processors only lower their price by an average of 26 cents per gram, suggesting that tax invariance has been violated in this setting.

²⁹We ignore the taxes applied to cultivators in this discussion because over 90 percent of the market was vertically integrated, and thus did not have to pay the tax and our price data for the remaining cultivators is not of high-enough quality to consider this part of the tax change directly.

We now consider the retail 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 percent and the sales tax switched from being on the tax-inclusive to the tax-exclusive price), the incidence calculations would be straightforward. The retail tax rate increased by 7 percent³⁰, and prices faced by consumers went up by 2.33 percent, so tax pass-through to the consumer would be 33 percent.

However, the processor tax also changed and this led to a decline in processor prices (i.e. retailers' marginal costs went down). We expect marginal cost changes will be passed through to consumers in the same way as tax changes, so our calculation of consumer incidence that ignores this processor price change is biased downwards. In fact, if the tax change was entirely revenue neutral and tax-invariance had held, we would expect no change at all in retail prices, not because there was no pass-through, but because there was no change in taxes to be passed-through. We see from Figure 10 that processor prices fell by 26 cents for an average gram of marijuana. This amounts to a 1.94 percent decrease from the equilibrium tax-inclusive price, \$13.49.³¹ Hence, the net change in taxes and marginal costs is 5.06 percent, so consumers bear 44 percent of the tax burden ($=.0233/.0506$).

We would be remiss to not mention the federal income tax in the discussion of incidence, though appreciating exactly how it changed in expectation in response to the tax change is challenging. The wording of the original Washington law suggested that retail firms may owe federal income taxes on the tax-inclusive price.³² This was part of the instigation for

³⁰The tax changes amount to a 7 percent 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 percent ($= \log((1 + .37 + .087)/(1.25 * 1.087))$).

³¹We calculate this as the log change, $0.0194 = \log((13.49 - .26)/13.49)$.

³²Here is the original language of the bill: "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

the reform of the tax system on July 1, 2015. An IRS memo was eventually written, which clarified that taxes were not owed on taxes paid to the state and this memo was retroactive to the previous year's taxes, 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 due to the change in retail and processor prices in response to the tax change. If they did not, the tax change on July 1, 2015 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. If firms believed that that the federal income tax always applied to the tax-exclusive price, we find that incidence is split approximately 42-68 between consumers and retailers – the only difference from the baseline incidence estimate is that firms can deduct processor prices from their federal income taxes and these went down in response to the tax change. As another benchmark, we would have found that taxes were completely passed through to consumers if 70 percent of firms believed there would be no federal income tax break due to the reform and 30 percent believed that there would be a tax break (or equivalently if all firms believed there was a 30 percent chance there would be a federal income tax break associated with the tax reform). For additional details on our incidence calculations that include the federal income tax, please see Appendix A.

apply.” Because of the last part of this sentence, retail firms were concerned/believed that they owed federal income taxes on their tax-inclusive price, too.

5 Conclusion

After the ballot measures passed during the 2016 elections are fully implemented, recreational marijuana will be legal for 21 percent 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 regulatory and taxation systems. Internationally, countries that have chosen to legalize marijuana, such as Canada and Uruguay, have had robust public debates about the design of the regulatory and taxation scheme in the face of black market availability.

We contribute to this debate by analyzing a unique natural experiment in Washington wherein the tax regime was reformed from a three-tiered structure of gross receipts taxes throughout the supply chain to a single excise tax at retail. We use novel, comprehensive administrative data on the universe of marijuana production and transactions within the state to understand how cultivators, processors, retailers, and consumers responded to the reform.

We find gross receipts taxes increased vertical integration, which decreased the efficiency of the market, given that non-vertically integrated transactions increased substantially, and 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 these gross receipts taxes may be larger than other taxes with similar collected amounts given their effects on decreasing the efficiency of production. This finding is particularly relevant given the diverse regulations present in the nascent marijuana markets across the United States. Colorado

requires vertical integration (retail firms must produce at least 80 percent of the marijuana they sell) and Washington allows some types of vertical integration (between cultivators and processors) and forbids integration between retailers with firms involved in marijuana's production. Our findings suggest that requiring vertical integration, as Colorado does, will decrease market efficiency.

While the previous literature has studied the elasticity of demand for marijuana within the black market, ours is the first to focus exclusively on the legal market for recreational marijuana. Some recent studies, notably Jacobi and Sovinsky (2016), have estimated the elasticity of marijuana to be -0.2 , our findings suggest demand is inelastic (-0.4) in the days surrounding 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 suggests the medium-run response to a price increase is elastic, which in turn suggests Washington is close to the peak of the Laffer curve. Given that the tax rates that were passed in California (15 percent), Maine (10 percent), Massachusetts (3.75 percent), and Nevada (15 percent), are considerably lower than Washington's rate, our findings suggest considerable state revenue may be left on the table.

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 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. Thus, the tax regimes

in Alaska, California, and Colorado may incentivize suppliers to produce higher priced and higher potency marijuana.

While our results suggest significant tax revenue may be left on the table in many jurisdictions, 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.

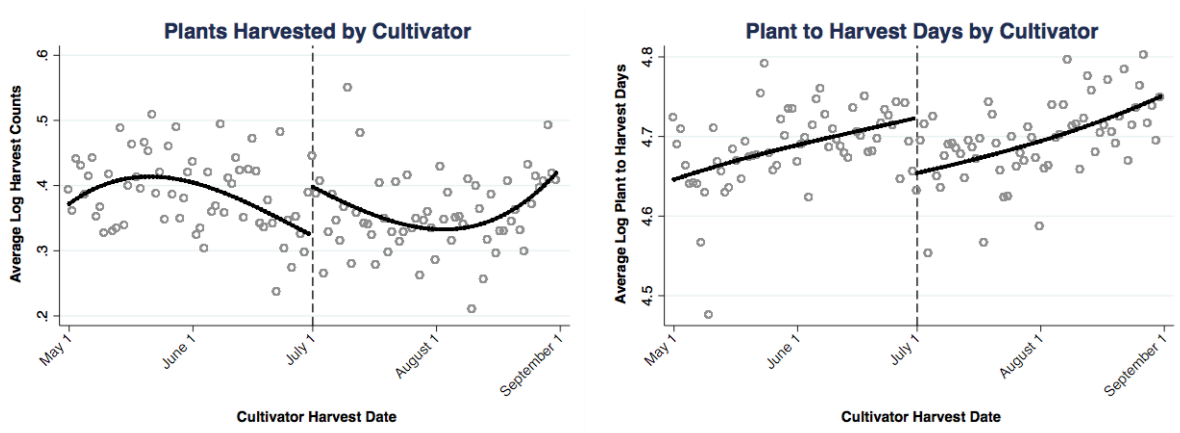
References

- [1] Anderson, D. Mark, Benjamin Hansen, and Daniel 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] Chetty, R., Looney, A. and Kroft, K (2009). "Salience and Taxation: Theory and Evidence." *American Economic Review*, 99(4), 1145-1177.
- [5] 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*.
- [6] Clements, K. W., & Zhao, X. (2009). *Economics and Marijuana: Consumption, Pricing and Legalization*. Cambridge University Press.
- [7] Cole, J. (2013). Memorandum for all United States attorneys. accessed from <https://www.justice.gov/iso/opa/resources/3052013829132756857467.pdf> on 07/20/2016
- [8] Dobkin, C., & Nicosia, N. (2009). "The war on drugs: methamphetamine, public health, and crime." *The American Economic Review*, 99(1), 324-349
- [9] 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.
- [10] 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.
- [11] 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.
- [12] Hall, W., & Lynskey, M. (2016). "Evaluating the Public Health Impacts of Legalizing Recreational Cannabis Use in the United States. " *Addiction*, 111: 1764-1773.

- [13] Jacobi, L. & M. Sovinsky. (2016). "Marijuana on Main Street? Estimating Demand in Markets with Limited Access." *American Economic Review*. 106(8): 2009-45.
- [14] Kenkel, D. S. (2005). "Are Alcohol Tax Hikes Fully Passed through to Prices? Evidence from Alaska." *The American Economic Review*, 95(2), 273-277.
- [15] 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.
- [16] Laffer, A. (2004). "The Laffer Curve: Past, Present, and Future."
- [17] La Corte, R. (2015) "Gov. Inslee signs recreational marijuana reform bill." *ap.org*. Associated Press, 30 Jun. 2015.
- [18] 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.
- [19] Miron, J., & Zwiebel, J. (1991). "Alcohol Consumption During Prohibition." *American Economic Review*, 81(2), 242-247.
- [20] Miron, J. A., & Zwiebel, J. (1995). "The Economic Case Against Drug Prohibition." *The Journal of Economic Perspectives* 9(4), 175-192.
- [21] Motel, S. (2015). Six Facts about Marijuana. accessed from <http://www.pewresearch.org/fact-tank/2015/04/14/6-facts-about-marijuana/> on 10/22/2016.
- [22] 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
- [23] Pogue, T. (2007). "The Gross Receipts Tax: A New Approach to Business Taxation," *National Tax Journal*, 60(4), 799-819.
- [24] Poterba, J. (1996). "Retail Price Reactions to Changes in State and Local Taxes." *National Tax Journal*, 49, pp. 165-176.
- [25] Testa, W. & R. Mattoon (2007). "Is There a Role for Gross Receipts Taxation?" *National Tax Journal*, 60(4), 821-840.
- [26] 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*.

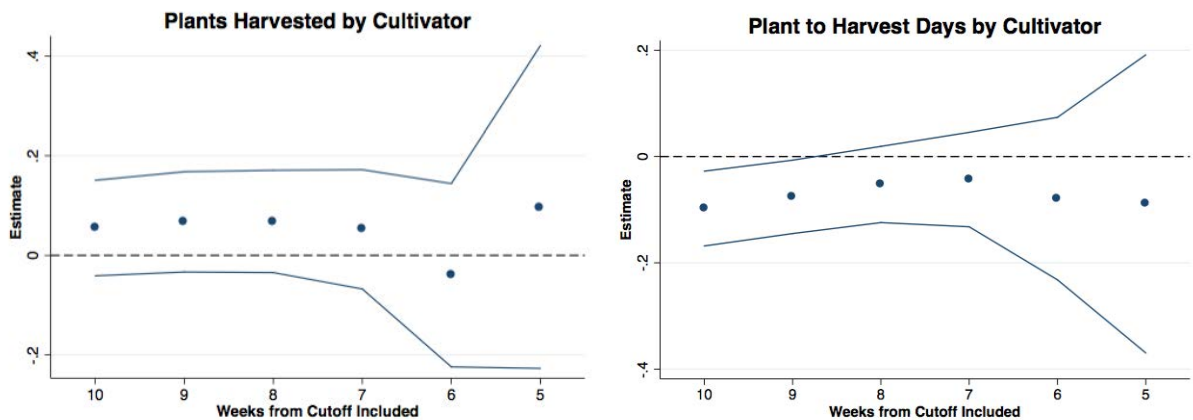
6 Figures and Tables

Figure 1: Cultivator Harvests and Days-to-Harvest



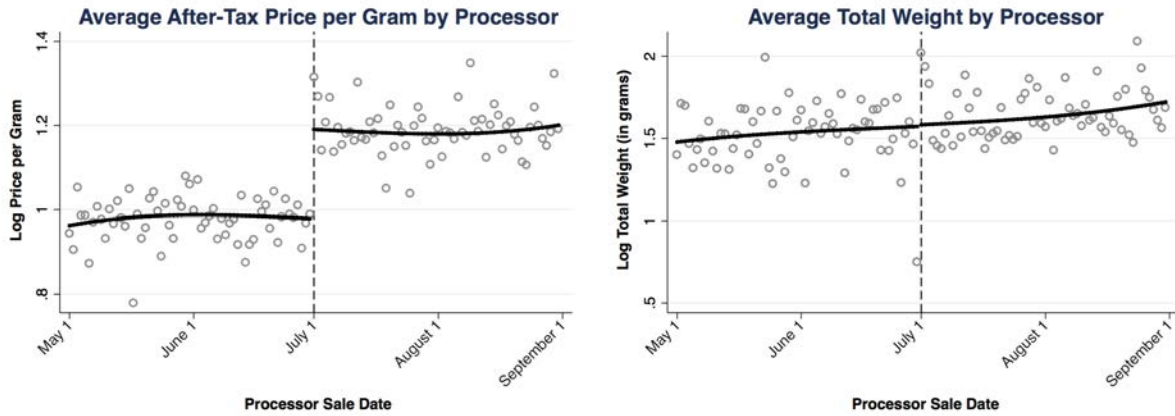
These figures are based on the estimates in Column (1) of Table 5. The solid line plots the estimated response to the tax change plus the polynomial. The scatterplot is the corresponding daily average of the dependent variable with estimates of the day-of-week, day-of-month, and cultivator fixed effects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



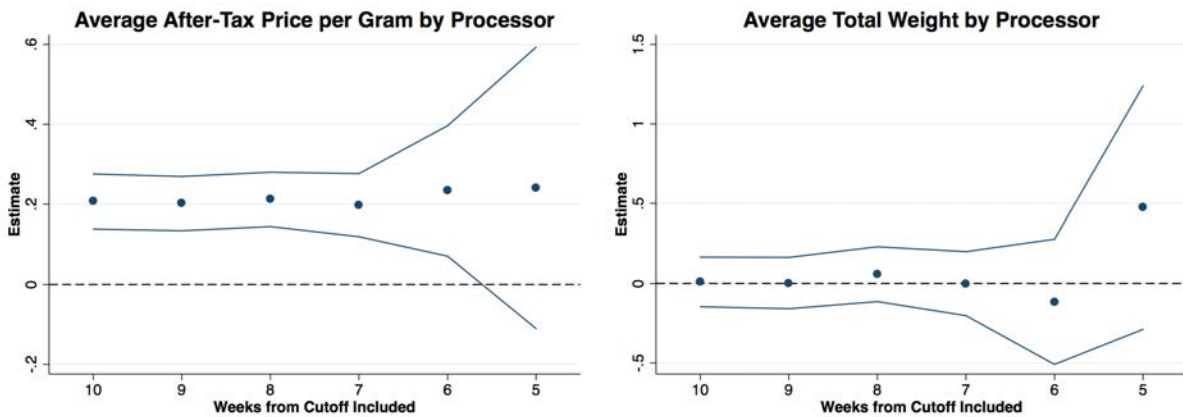
These figures consider the sensitivity of the estimates in Column (1) of Table 5 to the number of weeks of data we include on either side of the tax change. The dots mark the estimates for each bandwidth choice and the lines mark the 95 percent confidence intervals around these estimates.

Figure 3: Processor After-Tax Prices and Revenue



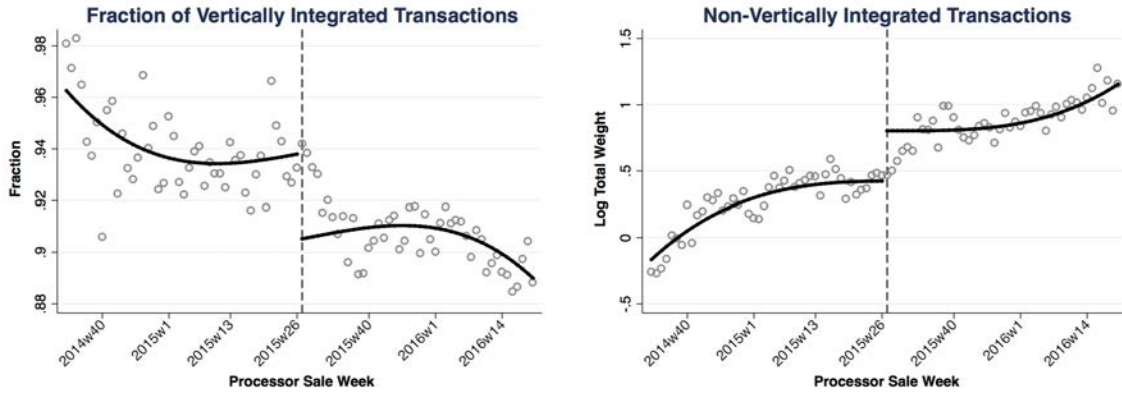
These figures are based on the estimates in Column (1) of Table 8. The solid line plots the estimated response to the tax change plus the polynomial. The scatterplot is the corresponding daily average of the dependent variable with estimates of the processor fixed effects removed. The vertical dashed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

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



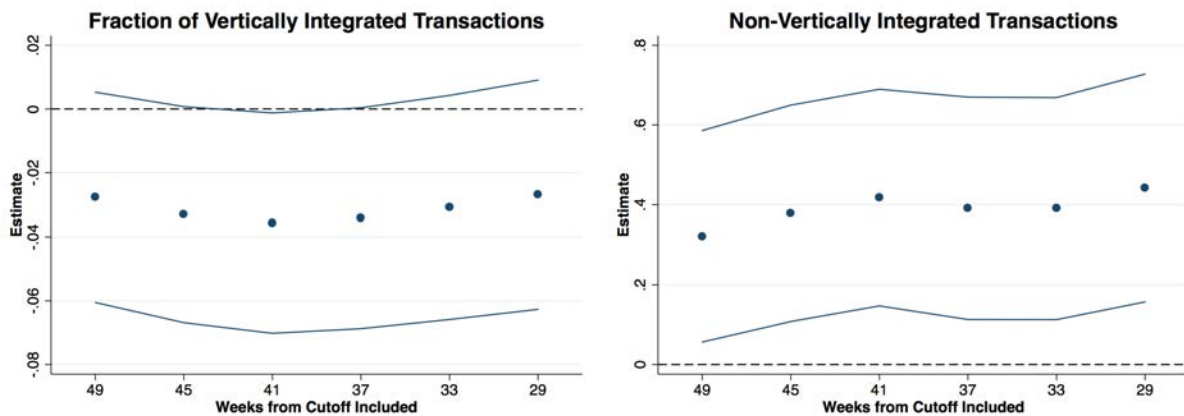
These figures consider the sensitivity of the estimates in Column (1) of Table 8 to the number of weeks of data we include on either side of the tax change. The dots mark the estimates for each bandwidth choice and the lines mark the 95 percent confidence intervals around these estimates.

Figure 5: Vertical Integration



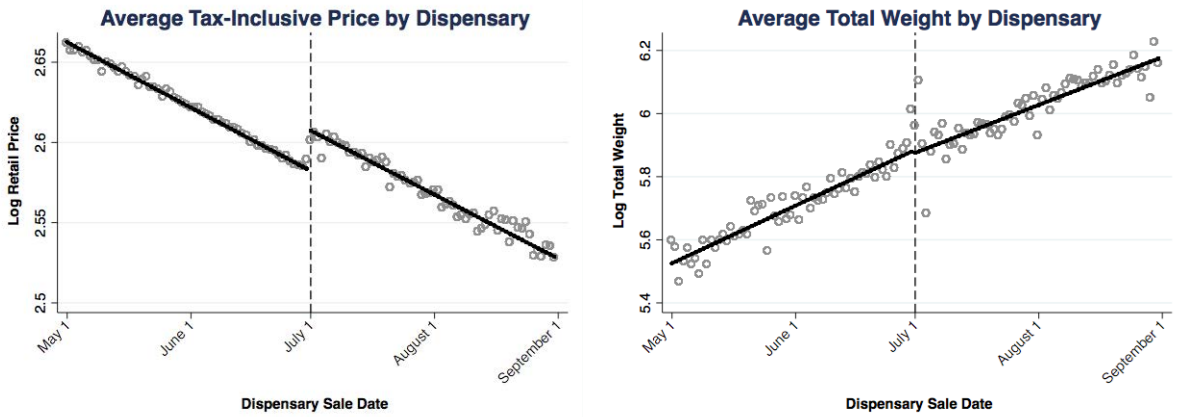
These figures are based on the estimates in Column (1) of Table 10. The solid line plots the estimated response to the tax change plus the polynomial. The scatterplot is the corresponding weekly average of the dependent variable with estimates of the processor fixed effects removed. The vertical dashed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

Figure 6: Vertical Integration Bandwidth Sensitivity



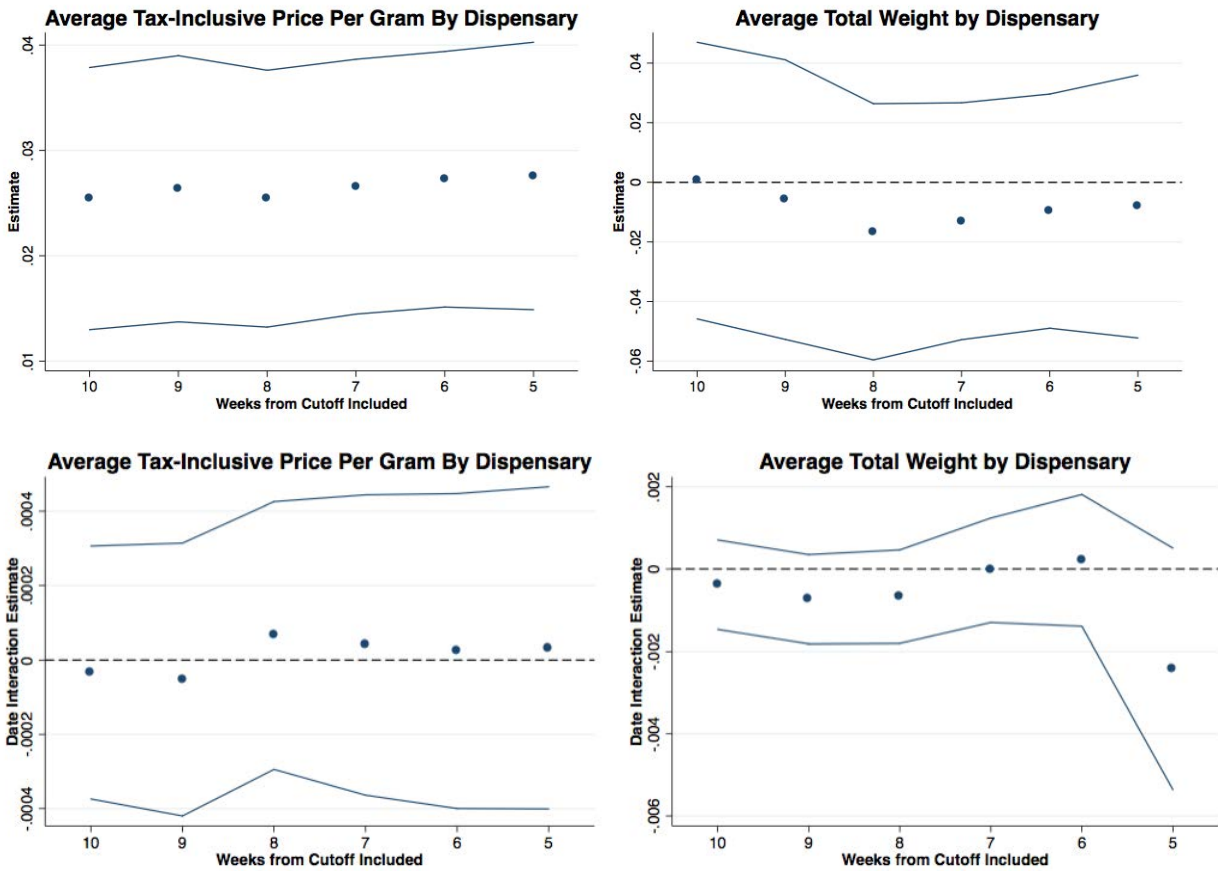
These figures consider the sensitivity of the estimates in Column (1) of Table 8 to the number of weeks of data we include on either side of the tax change. The dots mark the estimates for each bandwidth choice and the lines mark the 95 percent confidence intervals around these estimates.

Figure 7: Retail Tax-Inclusive Prices and Weight



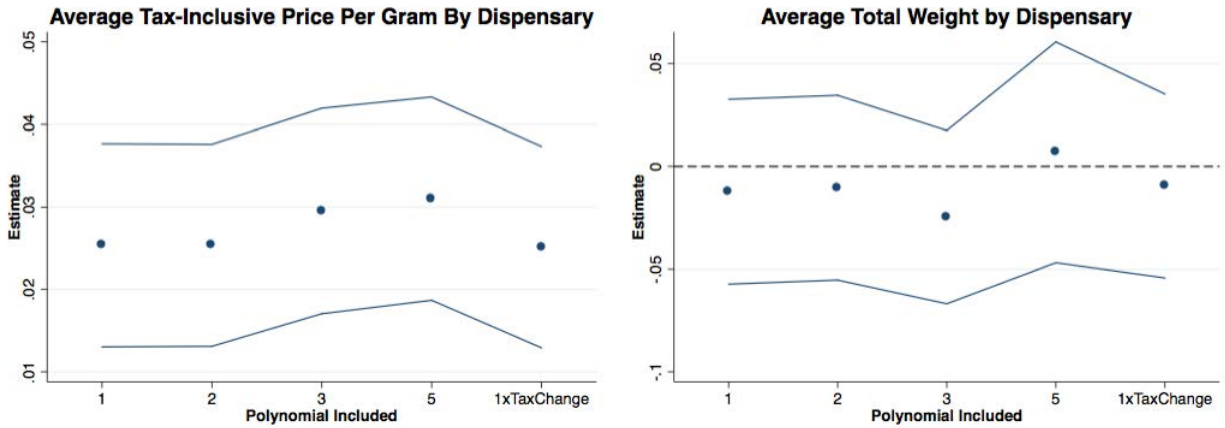
These figures are based on the estimates in Column (1) of Table 13. The solid line plots the estimated response to the tax change plus the date polynomial. The scatterplot is the corresponding daily average of the dependent variable with estimates of the day-of-week, day-of-month, and retail fixed effects 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



d 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 percent 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 percent 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.

Table 1: 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. Localities may impose additional retail taxes.
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	3.75%	Localities may impose additional 2% 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

Table 2: House Bill 2136 legislative history summary

Date	Activity
	Regular Session
February 17	Introduced in WA House, referred to committee
February - March	Passed by committees, substituted twice
April 10	Passed by House 67-28, referred to Senate
April 24	Accepted by Senate Committee, referred back to House, Regular Session Ends
	First Special Session
April 29	House passes 70-25, referred to Senate
May 1	Referred to Senate committee
May 28	Referred back to House, First Special Session Ends
	Second Special Session
May 29	Reintroduced in House, referred to committee
June 26	Removed from committee, passed by House 59-38
June 27	Passed by Senate 36-7
June 30	Signed by Governor

Source: <http://app.leg.wa.gov/bills/summary?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

Table 4: Cultivator Response to Tax Change

	(1) Plantings	(2) Harvests	(3) Days-to-Harvest
Tax Change	-0.005 (0.035)	0.075 (0.051)	-0.069* (0.036)
Observations	45,401	30,477	3,739
R-squared	0.106	0.104	0.569
Months Pre-Post	2	2	2
Polynomial Order	3rd	3rd	3rd
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, and cultivator location fixed effects. All dependent variables are in logs; 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 (0.051)	0.068 (0.057)	0.057 (0.045)	0.087* (0.048)	0.084 (0.066)	0.074 (0.051)
Days-to-Harvest	-0.069* (0.036)	-0.085** (0.037)	-0.058* (0.031)	-0.048 (0.032)	0.003 (0.050)	-0.074** (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 log of the tax-inclusive price per gram. The next row reports an estimate for the dependent variable log weight. 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, and cultivator location fixed effects. All dependent variables are in logs. 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
Fraction of Vertically Integrated Firms	6,514	0.93	0.22	0.93	0	1

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** (0.034)	0.212*** (0.034)	0.009 (0.085)	0.009 (0.030)	-0.009 (0.102)	0.054 (0.100)	-0.036 (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
Months Pre-Post	2	2	2	2	2	2	2
Polynomial Order	3rd	3rd	3rd	3rd	3rd	3rd	3rd
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, and processor location fixed effects. All dependent variables are in logs; 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*** (0.034)	0.209*** (0.035)	0.221*** (0.030)	0.238*** (0.032)	0.189*** (0.054)	0.213*** (0.034)
Weight	0.009 (0.085)	-0.018 (0.095)	0.047 (0.078)	0.014 (0.079)	0.129 (0.114)	0.006 (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
Day of Month Indicators	Yes	Yes	Yes	No	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 log of the tax-inclusive price per gram. The next row reports an estimate for the dependent variable log after-tax revenue. 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, 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** (0.018)	0.418*** (0.147)	0.047 (0.036)	-0.162** (0.081)	1.046** (0.517)
One Week Post	0.041*** (0.015)	-0.378*** (0.138)	-0.079** (0.037)	0.199*** (0.056)	-1.298*** (0.359)
Two Weeks Post	0.037** (0.016)	-0.342*** (0.122)	-0.007 (0.053)	0.195*** (0.054)	-1.271*** (0.347)
Three Weeks Post	0.031** (0.013)	-0.272*** (0.104)	-0.050 (0.036)	-0.059 (0.225)	0.083 (1.208)
Four Weeks Post	0.028** (0.011)	-0.192* (0.108)	0.018 (0.039)	0.186*** (0.050)	-1.215*** (0.322)
Five Weeks Post	0.012 (0.010)	-0.163* (0.094)	-0.063** (0.030)	0.182*** (0.048)	-1.188*** (0.310)
Six Weeks Post	0.017 (0.011)	-0.189* (0.097)	-0.048 (0.039)	0.178*** (0.046)	-1.160*** (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 month are included, but not reported. All dependent variables are in logs; the dependent variable is listed directly below the column number. NVWeight stands for 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** (0.018)	-0.038** (0.019)	-0.031* (0.017)	-0.040** (0.019)	-0.039* (0.021)	-0.034** (0.017)
Non-Vertical Weight	0.418*** (0.147)	0.427*** (0.154)	0.351*** (0.133)	0.406*** (0.156)	0.488*** (0.159)	0.407*** (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 non-vertical weight. Processor location fixed effects and a polynomial in the week are included. All dependent variables are in logs. 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

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*** (0.012)	0.021* (0.012)	0.007 (0.027)	0.022 (0.029)	-0.036 (0.030)	0.034 (0.030)	-0.018*** (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
Months Pre-Post	2	2	2	2	2	2	2
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 and day of the month indicator variables, 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*** (0.0066)	0.0228*** (0.0068)	0.0220*** (0.0062)	0.0237*** (0.0064)	0.0233*** (0.0066)	0.0233*** (0.0066)	0.0229*** (0.0066)	0.0241*** (0.0068)
Inventory Lot TI Price Interaction	-0.0001 (0.0002)	-0.0001 (0.0002)	-0.0000 (0.0002)	-0.0001 (0.0002)	-0.0001 (0.0002)	-0.0001 (0.0002)	-0.0001 (0.0002)	-0.0001 (0.0002)
Weight	-0.0095 (0.0228)	-0.0038 (0.0229)	-0.0079 (0.0207)	0.0135 (0.0215)	-0.0085 (0.0220)	-0.0092 (0.0222)	-0.0096 (0.0228)	-0.0072 (0.0228)
Weight Interaction	-0.0010* (0.0006)	-0.0008 (0.0006)	-0.0010* (0.0005)	-0.0005 (0.0006)	-0.0009 (0.0006)	-0.0009 (0.0006)	-0.0010* (0.0006)	-0.0010* (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 reports the estimate on the Tax Change indicator variable. The second row reports the estimate on the interaction between the running variable and the Tax Change. The next two rows report the same estimates for the dependent variable log weight. The elasticity of demand implied by these estimates is calculated by dividing the weight estimate by the price estimates in the second 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 first and third rows), day indicator variables for June 30 - July 5, day of the week and day of the month indicator variables, retail location fixed effects, and a linear running variable/polynomial in date. All dependent variables are in logs. Standard errors clustered by retailer are in parentheses.

Appendices

A Federal Income Tax Incidence Calculations

In this appendix, we consider how our incidence results would be affected if we incorporated the federal income tax into our calculations. Without considering the federal income tax, we found that consumers bore about 44 percent 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 and, as we described in the main text, depends on whether firms believed that they always owed taxes on the tax-exclusive price of marijuana or whether they believed that they owed taxes on the tax-inclusive price before the tax reform and on the tax-exclusive price after the tax reform. While we can't say exactly what incidence looks like when we incorporate federal 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.³³

We don't 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 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) their only income source is their marijuana business, (2) their

³³We 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 don't see, so our calculations of the federal income tax liability would be much less accurate.

only cost of goods sold is the marijuana purchased, (3) they do not itemize, and (4) 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 percent increase from the equilibrium tax-inclusive price, \$13.49.³⁴ Hence, the net change in taxes and marginal costs is 5.50 percent, so consumers bear 42 percent of the tax burden ($=.0233/.0550$).

As an alternative benchmark, we consider what happens when 70 percent of firms believe that federal income taxes are always assessed on the tax-exclusive price and 30 percent 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 ($=.3*\$2.74+0.7*\1.54) on average before the tax reform. Federal income taxes on a dollar of marijuana declined by 36 cents, which amounts to a 2.7 percent increase from the equilibrium tax-inclusive price, \$13.49.³⁵ Hence, the net change in taxes and marginal costs is 2.33 percent, so consumers bear 100 percent of the tax burden ($=.0233/.0233$). As the fraction of firms that expected to pay taxes on the tax-inclusive price before the tax change increases above 30 percent, pass-through to consumers would be more than 100 percent. And, as it decreases below 30 percent, pass-through to consumers would be less than 100 percent.

³⁴We calculate this as the log change, $0.0044 = \log((13.49 + .06)/13.49)$.

³⁵We calculate this as the log change, $0.027 = \log((13.49 - 0.36)/13.49)$.