

NBER WORKING PAPER SERIES

VERTICAL INTEGRATION AND PRODUCTION INEFFICIENCY
IN THE PRESENCE OF A GROSS RECEIPTS TAX

Benjamin Hansen
Keaton S. Miller
Caroline Weber

Working Paper 28478
<http://www.nber.org/papers/w28478>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2021

The authors would like to thank David Agrawal, J.S. Butler, Estelle Dauchy, Naomi Feldman, Enda Hargaden, Austin Nichols, Nathan Seegert, Monica Singhal, Jeffrey Smith, Alex Yuskavage and seminar or conference participants at University of Wisconsin-Madison, University of Kentucky, ABFM, NTA, WEAI and the Utah Tax Invitational (U-TAXI) for helpful comments. This paper previously circulated as: “The Welfare Costs of a Gross Receipts Tax”. 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Benjamin Hansen, Keaton S. 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.

Vertical Integration and Production Inefficiency in the Presence of a Gross Receipts Tax
Benjamin Hansen, Keaton S. Miller, and Caroline Weber
NBER Working Paper No. 28478
February 2021
JEL No. H21,H25,H26,H30,H71,I28,L51,L6

ABSTRACT

We quantify the effects of a gross receipts tax (GRT) on vertical integration for the first time. We use data from the Washington state recreational cannabis industry, which has numerous advantages including a clean natural experiment: a 25% GRT imposed on cannabis firms was subsequently replaced by an excise tax at retail. We find the short-run elasticity of vertical integration with respect to the intermediate good net-of-tax rate is -0.15 and the long-run elasticity is more than twice as large. We find these incentives lead to large output losses – production increases by 23 percent when the GRT is eliminated.

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

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

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

1 Introduction

Gross receipts taxes (GRTs) or turnover taxes are taxes on firm revenue with no deductions for costs. When imposed on industries with dis-integrated supply chains, GRTs create tax pyramids or tax cascading—final goods are taxed multiple times throughout production. This pyramiding incentivizes firms to vertically integrate to avoid layers of intermediate good taxation (a contracting friction in the sense of Coase, 1937) even when such integration is inefficient from a production cost perspective (Williamson, 1971) and would not lead to strategic gains (see e.g. Salinger, 1988). Commodity or income taxes are therefore preferred in environments without enforcement concerns (Diamond & Mirrlees, 1971). In environments with limited enforcement capacity, however, GRTs can be preferred because they reduce evasion relative to a VAT (Best et al., 2015).

Despite this theoretical understanding of the tradeoffs involved in imposing GRTs, it has been difficult for economists to study these policies empirically due to the need for production data covering the entire supply chain for a sizeable product combined with a plausible natural experiment. Empirical analyses of GRTs are important in part due to the re-emergence of this policy instrument in the United States – nine U.S. states have GRTs (see Section 2.1) – and around the world (Ernst & Young, 2020). At the same time, while production efficiencies are one of the primary features of theories of vertical integration and play a key role in antitrust policy (Blair & Kaserman, 2014), it has been difficult to systematically measure the effects of integration decisions due in part to selection concerns and difficulties finding appropriate sources of variation (Bresnahan & Levin, 2012).

We provide the first empirical evidence of the vertical integration incentives and resulting

production inefficiencies created by GRTs by examining the recreational cannabis market in Washington state. This setting is ideal—it combines data that spans the universe of intermediate and final goods produced by a multi-stage supply chain with a reform that removed a GRT and replaced it with an approximately revenue-neutral sales tax. By comparing firm behavior before and after the reform, we estimate the elasticities of vertical integration and production with respect to the gross receipts net-of-tax rate. We thus complement the work of Best et al. (2015), who provide evidence on the evasion elasticity in a setting with limited enforcement but are unable to examine vertical integration or production inefficiencies.

We begin in Section 2 by describing the institutional features of Washington’s cannabis industry. Washington’s regulations created a supply chain with three types of firms: cultivators, which grow cannabis plants, manufacturers, which transform raw plant material into final products, and retailers, who sell final products to end-users. Crucially, vertical integration between cultivators and manufacturers is allowed but not required.¹ To comply with state law, firms must provide detailed information about their operations to the state government. These unique “seed-to-sale” administrative records, which we describe in Section 3, track the entire legal supply chain (cannabis cannot legally be sold across state lines) and are verified through frequent government audits.

Prior to July 1, 2015, a 25% GRT was assessed on each firm within Washington’s supply chain; thus a firm which was comprised of a vertically integrated cultivator and manufacturer enjoyed a potential cost advantage over dis-integrated competitors. On July 1, the GRT was replaced by a 37% excise tax at retail. This change was unexpected by market participants: the reform was passed during a special session of the Washington Legislature on June 27,

¹Vertical integration between retailers and any other type of firm is prohibited.

2015, and signed by the Governor on June 30. In Section 4 we measure the effects of this reform using an interrupted time series regression; that is, we ask how industry outcomes change in the weeks after the reform (excluding an adjustment period) relative to pre-reform outcomes. Identification rests on the assumption that, after controlling for product characteristics, industry outcomes would not have changed (relative to a baseline trend) in the absence of the reform. We perform placebo analyses and additional robustness checks which provide no evidence to reject this assumption.²

We find the GRT encouraged vertical integration: the share of vertically-integrated cannabis (defined as product for which the cultivator and manufacturer were the same firm) fell by 4.6 percentage points immediately following the elimination of the GRT. To understand the long-run effects of a GRT, we compare post-reform entrants (who presumably had not paid fixed costs at the time of the reform) to incumbents: we find the difference in the relative share of vertically-integrated cannabis across the two groups is more than twice as large as the short-run effects within the incumbent group. We estimate that the long-run elasticity of vertical integration with respect to the intermediate good net-of-tax rate is -0.39.

We find the vertical integration incentives created by the GRT led to substantially lower production (as measured by the total quantity of final goods produced); we estimate a production elasticity with respect to the intermediate good net-of-tax rate of 0.68. The increase in production, as expected, is driven by the cultivators which participate in the non-vertical market (i.e. those cultivators which sell intermediate goods to manufacturers

²We employ this approach rather than a state difference-in-differences design as the only potential comparison state is Colorado, which had a different regulatory and industry structure, and for which equivalent data is unavailable. In particular, Colorado required a level of vertical integration between production and retail. The assumption that outcomes in the two states co-move in the period of the reform is likely much stronger than the assumptions we impose.

with which the cultivators in question are not vertically integrated) post-reform. We provide evidence that these estimates are not biased by tax evasion or inaccurate reporting, and are not driven by reverse causality through the price channel (Alfaro et al., 2016; McGowan, 2017).

Lastly, we explore heterogeneity in the responses to the reform across firms with different production capacities. We find that manufacturers with little-to-no production capacity (as determined by the size constraint of their cultivation license, if present) bear the brunt of the GRT. When the GRT is eliminated, these firms increase their purchases of non-vertically-integrated intermediate goods which allows them to increase their sales by approximately 80%. Cultivators across the capacity distribution increase their production by approximately 20%. We conclude by discussing the policy implications of these findings in Section 5.

More generally, our results speak to the effects of transitions from high-cascade (or high-pyramiding) taxes to low-cascading taxes. The most common transition studied is from a retail sales tax (RST) or goods and services tax (GST) to a VAT. While RSTs are generally imposed with the intent to avoid taxing business inputs, estimates suggest that more than 40 percent of business inputs are taxed in practice and thus pyramiding is a concern (Wildasin, 2001; Ring, 1999; Smart & Bird, 2009; Phillips & Ibaid, 2019). The taxation of business inputs is expected to be much lower under a VAT (Kopczuk & Slemrod, 2006). For example, (Smart & Bird, 2009) study a transition from a GST to a VAT in Canada and argue that capital investment increased largely due to a decrease in tax pyramiding. In our setting investment goods are not taxed due to the industry-specific nature of the tax; our estimates thus represent a lower bound on the total inefficiencies created by the pyramiding structure of the GRT. Agrawal & Zimmerman (2019) study a transition from a first-point-of-sale retail

tax to a VAT in India and find a substantial increase in retail sales. Relative to this literature, we are the first to identify the consequences of a cascading tax on vertical integration and the resulting production inefficiencies it creates.

Our work also contributes to the literature examining the empirical effects of vertical integration decisions.³ Much of this work focuses on disentangling the strategic (i.e. foreclosure) and efficiency motives for integration (see e.g. Crawford et al., 2018). In our setting, the degree of competition at each stage of the supply chain implies that strategic incentives are unlikely to drive integration decisions and thus our results are primarily informative about productive efficiency (Hortaçsu & Syverson, 2007). We are unique in studying a policy-induced friction that leads to *too much* vertical integration rather than (for example) frictions driven by financial markets (Acemoglu et al., 2009), technology differences (Acemoglu et al., 2010), or quality differences (Hansman et al., 2020).

2 Background

2.1 Gross Receipts Taxes in the United States

GRTs, once in decline, have made a recent comeback (Mikesell, 2007a,b; Pogue, 2007; Testa & Mattoon, 2007; Kaeding, 2017), with nine U.S. states now imposing a GRT (see Table 1).⁴ The GRT rate is typically low – often less than 1 percent. However, Hawaii, New Mexico, and Washington all charge some rates approaching those of typical sales taxes. Most U.S. state GRTs are ‘impure’ in the sense that they do not tax all activity in all sectors at all

³See (Bresnahan & Levin, 2012) for a review.

⁴For a discussion of GRTs outside of the U.S., see Best et al. (2015).

stages of production at the same rate without allowing any deductions for costs.⁵

The re-emergence of the GRT in the United States is likely partly because the large tax base means a low rate can generate large revenues, partly because the GRT taxes services and RSTs generally do not (and other tax-base related issues), partly because RSTs already tax a substantial fraction of business inputs in practice (Phillips & Ibaid, 2019; Ring, 1999; Smart & Bird, 2009; Wildasin, 2001), and partly because the GRT tax base provides limited opportunities for firms to engage in tax evasion relative to profit-based taxes like the VAT and corporate income tax (Best et al., 2015). This last issue is likely important in environments with limited enforcement capacity, such as some developing countries (Carrillo et al., 2017) or in markets with many small businesses where enforcement can be challenging even in developed countries (Slemrod et al., 2017).

However imposing a GRT (or any cascading tax) is not without potentially significant costs; tax pyramiding, in theory, leads to inefficient vertical integration and production inefficiencies (Diamond & Mirrlees, 1971; Best et al., 2015). Even accounting for some of the practical issues discussed in the previous paragraph, a GRT is not expected to be a second-best tax (at least not without significant modification) (Pogue, 2007; Testa & Mattoon, 2007) except in the case of substantial VAT tax evasion where the VAT is the other alternative feasible tax instrument (Best et al., 2015). While the production inefficiencies from inefficient vertical integration created by the GRT (or other cascading taxes) are theoretically

⁵For example, Hawaii and Washington both have separate rates for retail and wholesale. In Hawaii, this may be intended to mitigate tax pyramiding as the wholesale rate is 0.5% while the retail rate is 4%. This does not appear to be the goal for Washington as their wholesale rate is slightly higher than their resale rate. New Mexico has many special rates to address this as well. Delaware, Nevada, Texas, and Washington vary their tax rate by industry to mitigate the otherwise large inequities across sectors with different profit margins. Partially addressing the same issue, Oregon and New Mexico exclude grocery stores. Most states provide a standard deduction, and two states – Oregon and Texas– allow firms to deduct some costs.

understood, their magnitude has never been estimated – this is the goal of the present paper. If a GRT is chosen because of limited tax enforcement capacity, the estimates in our paper reflect the real costs of poor enforcement. The other major drawback of a GRT is that it is expected to be inequitable by construction, creating highly unequal tax burdens across sectors with different profit margins (Mikesell, 2007a,b; Testa & Mattoon, 2007; McClure, 2017).

2.2 Washington’s Cannabis Industry

The setting for our empirical analyses is the adult-use cannabis market in Washington State, which opened in July 2014 after cannabis was legalized by ballot initiative in November 2012. We have written elsewhere about the history of this market (Miller & Seo, 2018; Hansen et al., 2020). Here we describe the features of the market that are key to our analysis and discuss ways in which we expect insights from this market will map onto other industries.

Washington’s legal structure created three types of firms: cultivators grow and harvest cannabis plants, manufacturers transform harvested plant material into cannabis products and sell them wholesale to retailers, and retailers sell final goods to consumers.⁶ Potential entrants have to pass background checks and undergo a lengthy regulatory process requiring substantial capital investment before entry. These rules impose substantial burdens on potential entrants that suggest that those with substantial financial and human capital are far more likely to successfully enter the market.

By law, cannabis firms may possess both a cultivation and a manufacturing license.

⁶Washington law refers to cultivators as ”producers” and manufacturers as ”processors”. We choose nomenclature to represent functional equivalents across a broad variety of industries with intermediate goods.

However, any owner of a retail license may not have any ownership interest in any other type of license, and vice versa.⁷ Thus, while vertical integration between cultivators and manufacturers is possible and indeed common, retailers are independent.

Cultivator licenses come with capacity constraints—there are restrictions on the square footage that can be used for “live plant production” (WAC 314-55-010) and firms may not merge to increase capacity. Licenses are divided into three tiers. A Tier 1, 2, or 3 licensee may have up to 2,000, 10,000, or 30,000 square feet of plant production, respectively. These constraints were set by the regulator before any licenses were issued. To our knowledge, they were set without analysis of the efficient scale of production. For convenience, we refer to a manufacturer without a cultivation license as a “Tier 0” firm.

The cultivation of cannabis is capital-intensive relative to other agricultural products, owing to the way in which cannabis plants generate differing levels of the psychoactive substances Δ 9-tetrahydrocannabinol (THC) and cannabidiol (CBD) in response to growing conditions (Caulkins, 2010; Aizpurua-Olaizola et al., 2016). Cultivators may grow cannabis plants outdoors, in a controlled indoor facility, or using a combination of indoor and outdoor spaces. The average time from initial planting to full harvest maturity is approximately 4 months, depending on the reproduction method—plants grown from cloned seedlings require additional capital and labor inputs but mature slightly faster than plants grown from seed. Cultivators may also delay the maturation process by controlling the available light.

The predominant final good by weight and sales volume is “usable marijuana”, which consists of flowers of the cannabis plant which have been dried, cured, and packaged into sealed packets containing a set weight (e.g. 1 gram, 3.5 grams, or 7 grams). Our empirical

⁷These restrictions extend to relatives of licensees.

analyses focus on these “usable marijuana” products exclusively, which is about 75% of the total market.⁸ Usable marijuana products are differentiated by the “strain” of the plant— analogous to different cultivars of e.g. apples or grapes—and the potency of the flowers as measured by the concentration of THC and CBD (Amin & Ali, 2019).

The manufacture of “usable marijuana” products from raw plant material is also time- and capital-intensive—it takes about six weeks, on average. Plant material must be dried in a temperature- and humidity-controlled location for over a week, “trimmed” of non-flower material (leaves and stems), and then cured over multiple weeks in a lower temperature and humidity environment (Green et al., 2001). The cured flowers are then divided into packages of equal weight for retail sale. Given these space and climate-control requirements, manufacturers may not easily increase throughput without significant capital investment.

The market features a closed supply: all cannabis sold by retailers is grown in the state, and every ounce grown legally within the state is sold at a Washington retailer. These rules are enforced through the state’s “seed-to-sale” traceability system, which tracks each plant from cultivation through manufacturing and retail. This system was implemented to respond to the informal federal regulations created in response to the legalization efforts in Washington and Colorado (Cole, 2013). The system provides information that can be used to check for tax evasion: retailers cannot sell cannabis without manufacturing records, which generally forces manufacturers to report accurately.⁹ Reporting is enforced through frequent

⁸Due to limitations of the traceability system, the “usable marijuana” category we consider contains two types of products: both raw dried flowers and pre-rolled joints, which include some small additional value. The data on the rest of the market is mostly unusable because while we observe revenue from each sale, we do not have reliable quantity information (we observe the number of units sold of each product, but a unit of one product might contain many more doses than a unit of a different product) which is important for capturing changes in activity separately from compositional changes.

⁹Retailers can under-report their sales, but such behavior is detectable as retail sales can be compared to purchases from manufacturers. Our estimates are unaffected by dropping the few retailers that engage in

in-person audits—cultivators and manufacturers face an average of four in-person visits from auditors per year—backed by civil and criminal penalties for non-compliance.

To summarize: the cultivation and manufacturing steps in the cannabis production process involve substantial fixed costs that vary with the firm’s desired capacity and which do not functionally overlap.¹⁰ Converting capital from one use to another is costly—grow rooms require different infrastructure than drying and curing rooms. Expanding capacity requires investments on par with those incurred at entry—in other words an incumbent firm seeking to increase their processing throughput by, e.g. 100 kilograms of plant material per month, faces expansion costs that are similar to the incremental costs faced by a potential entrant considering the addition of 100 kilograms of plant material per month to their planned entry capacity. Indeed, insofar that potential entrants may be less constrained in their physical location choice, the marginal fixed cost of additional capacity may be lower for new entrants than for existing firms. Finally, the differing capital requirements suggest that the cultivation and manufacturing steps may operate most efficiently at different scales – in other words there is no reason *a priori* to suggest that vertical integration is efficient from a production process perspective, though it may be profit-maximizing when considering competitive effects (Bresnahan & Levin, 2012). We return to this point in our empirical analyses.

This qualitative description of the industry corresponds to accounts from industry participants. Prior to the elimination of the GRT, cultivators reported that they did not want to invest in the equipment to be manufacturers too, but could not afford to do otherwise. Moreover, even if a firm was both a cultivator and a manufacturer, it would have still liked

significant under-reporting.

¹⁰Atalay et al. (2014) argue that vertical integration in many contexts is driven by the desire for efficient intrafirm transfers of intangible inputs. While it is possible that intangible inputs play a role in the cannabis production process, the reform we study is unlikely to affect these inputs.

to sell some of its raw material to other manufacturers or to have been able to purchase from other cultivators, but both options were often made too unattractive by the existence of the 25% tax. These accounts are largely borne out in the data (see Section 4.1).

The extent to which our analysis is externally valid depends in part on the degree to which other industries face a tradeoff between strategic and efficiency incentives that mirror the tradeoff in Washington’s cannabis industry. One immediate comparison is the wine industry. Wine is produced in a similar process: a lengthy, capital-intensive cultivation step followed by a lengthy processing step requiring non-overlapping capital investments. Regulations generally impose a divide between producers and retailers to avoid “tied-houses” (Corsinovi & Gaeta, 2019). The wine industry features both vertically integrated production (“estate-grown” products) and vertically dis-integrated production. Wine producers at all scales generally combine grapes grown internally with those grown by specialized vineyards in order to produce an arrangement of end-consumer goods; capacity constraints generally bind for at least one stage of the production process (Allen & Lueck, 2019).

The tradeoff between strategy and efficiency is not limited to agriculture; all that is required is a multi-stage production process where capacity constraints bind differently throughout the stages. These features are present in numerous manufacturing and service industries. For example, smartphones can be thought of as bundles of components (CPU, memory, screen, cellular transmitter, operating system) packaged by an integrator and sold by a financially independent retailer (Fan & Yang, 2020). Each intermediate good is produced in a process requiring independent capital at different scales. Alternatively, Kikuchi et al. (2018) present a model where coordination costs, transaction costs, and diminishing returns to management lead ex-ante identical entrepreneurs to organize into a production

chain with multiple steps which would be affected by a pyramiding tax in a similar way.

Finally, we note that our focus is on the vertical integration incentives between cultivators and manufacturers; as retailers were not allowed to vertically integrate, we do not analyze their behavior in this paper. We thus examine the effects of a GRT on vertical integration and production inefficiency when there is a single supply chain link. The inefficiencies created by a GRT will increase with additional links, so the estimates in this paper are a lower bound.

2.3 Washington’s Gross Receipts Tax and Its Reform

Washington’s initial cannabis tax regime consisted of a 25% GRT applied to each firm within the cannabis industry. Thus, cultivators remitted the tax when they earned revenue from selling plant material to manufacturers, manufacturers remitted the tax when they earned revenue from selling final goods to retailers, and retailers remitted the tax when they sold final goods to consumers. The retail tax was required to be included in the posted price making it functionally equivalent to other excise and sales taxes. Vertically-integrated manufacturers did not have to pay taxes on intra-firm transfers of cannabis from the ‘cultivation side’ to the ‘manufacturing side’.

The reform we analyze eliminated the 25% taxes on cultivators and manufacturers and increased the retail excise tax from 25% to 37% while modifying the tax base. This change was designed to be revenue-neutral—in other work we find that it was slightly revenue-decreasing on a per-unit-sold basis (Hansen et al., 2020). Other regulations relevant to the market operation described above were largely unaffected. One exception was the devolution of additional zoning powers to local jurisdictions—our baseline analyses hold the set of firms

fixed before and after the reform to account for any downstream impacts of this change.

Our identification strategy assumes that the policy change was unanticipated by market participants. While this is ultimately an empirical question that we examine below, it is helpful to summarize the bill’s history. The bill originated in and passed Washington’s House (as H.B. 2136) midway through the 2015 Regular Session. While a Senate committee recommended passage on the last day of the session, the full Senate declined to consider the bill. A similar pattern occurred when the bill was reintroduced in the First Special Session: the House quickly passed the bill, and the Senate chose not to take action. Finally, at the very end of the Second Special Session, June 27, the bill received a full Senate vote. The Governor signed it on June 30, and the law went into effect the next day. Contemporaneous media reporting portrayed the industry as unprepared for the change, with one retail store manager quoted as follows: “This is supposed to happen tomorrow. You have a few hours to change an entire market’s pricing structure. It is an exceptionally short window for such a tremendous change” (LaCorte, 2015).

3 Data and Methods

Our data consist of administrative records from the “traceability” (or seed-to-sale) system maintained by the Washington State Liquor and Cannabis Board (WSLCB). We obtain data on all plants, products, and sales. Firms and locations are given unique identifiers. Each plant is registered at the time of planting. Firms record the provenance of the plant material (e.g., a clone or a seed) as well as the strain.¹¹ Once harvested, flowers and other

¹¹Strains are defined by the cultivator.

plant material are collected and converted into a new “inventory lot” that is assigned a unique identifier (ID); products or material within a single inventory lot are assumed to be homogeneous. These intermediate products may progress through several processing steps before wholesale distribution.

The last processing step is the division of a large wholesale inventory lot of final goods described above into multiple smaller inventory lots with unique IDs for sale to individual retail stores. When lots are sold to retailers, the tracking system records the date, the IDs involved (and thus the quantity), and the price of the transaction. Consequently, an inventory lot ID uniquely identifies the retailer, manufacturer, and cultivator, as well as the strain and package size.¹² We observe each wholesale and retail sale and link the price, quantity, and transaction times to the relevant inventory lots.

We define an inventory lot of cannabis as “vertically integrated” cannabis if it was cultivated and manufactured by the same firm. We denote an inventory lot as “vertically dis-integrated” or “non-vertically integrated” cannabis if it was cultivated and manufactured by different firms.¹³ Vertically dis-integrated activity originates from two sources: (1) a firm is vertically dis-integrated (so a cultivator but not a manufacturer or vice versa) – for these firms, every transaction will be vertically dis-integrated. This is fairly rare in our setting. (2) a firm is vertically integrated, but the cultivator chooses to sell some cannabis to another manufacturer or vice versa. This is more common in our setting.

Table 2 summarizes these data by firm-week and cultivator license tier across the 16 weeks prior to the reform date. For cultivators, we restrict to indoor-only firms to avoid

¹²A small number of lots have multiple package sizes, which we identify and correct for.

¹³Capacity constraints or ebbs and flows in production may lead a cultivator which is part of a vertically integrated firm to contract with a different manufacturer which is part of a separate vertically integrated firm.

the seasonality of outdoor plantings. While plantings and sales may vary from week-to-week, on average firms engage in more than one transaction per day. The largest indoor cultivators planted an average of 185 plants per week, whereas the smallest planted an average of 13 plants per week, roughly in line with the capacity constraints imposed by the tiered cultivation licensing system. On the manufacturing side, the largest firms sold an average of 2.7 kg of finished product per week in 25 transactions with retailers, whereas the smallest firms sold an average of 0.5 kg in 5 transactions. Across sizes, the share of vertically integrated cannabis was similar at approximately 95%.

We merge this data with two additional firm-level datasets obtained from the WSLCB. The first contains the dates of all audits conducted and any violations given during those audits. We create a tax evasion (or reporting inaccuracy) indicator for each firm which is equal to one if that firm received a tax or traceability reporting violation in the six months before or after the reform. The second dataset contains indicators for whether each cultivator grows their cannabis indoors, outdoors, in a greenhouse, or some combination of these.¹⁴

To study the effects of GRTs on vertical integration, we focus on the fraction of manufacturer-retailer cannabis transactions that are not vertically integrated as our outcome of interest.¹⁵ We do not expect any change driven by the reform to be immediate because, on average, it takes six weeks for manufacturers to process intermediate goods into final goods (as measured by the average time between purchasing plant material from cultivators and selling

¹⁴Unfortunately, the spreadsheet provided by WSLCB is incomplete; among other things, it only includes firms operating as of July 2019. As a result, we are only able to match 82 percent of the firms in operation at the time of the reform to this data set. In our analysis, we sometimes look exclusively at indoor firms, and it's possible with this restriction we drop some firms that we are unable to identify. Firms missing from this data set plant at a similar frequency and volume to firms we are able to observe and our estimates are statistically indistinguishable if we include all missing firms in our regressions as well.

¹⁵In Appendix Figure A.1, we consider another outcome: the volume of non-vertically integrated cannabis.

usable marijuana produced from that material to retailers). We therefore treat the six weeks immediately following the reform as an adjustment period to isolate the true effect of the reform. In Appendix A we document an immediate increase in firms participating in the non-vertical market for plant material.

To examine the effects of a GRT on production, we examine cultivator plantings because planting is the first step in the production process. It is feasible for cultivators to respond to the reform by increasing their plantings immediately, so we do not exclude any period of time after the reform.¹⁶

We aggregate the data by firm location-week and perform a few minor data cleaning steps.¹⁷ We keep all manufacturers that sell in more than one week during the time span of our analysis. In the planting market, to focus on actual production, we restrict our analysis to plants which were subsequently harvested (plants may not be harvested because, for examples, they died, were thinned out, were converted to clones for subsequent plantings, or were the male of a hybrid).¹⁸ We exclusively examine the indoor-only cultivator market to eliminate cyclical effects.¹⁹

We analyze these outcomes with an interrupted time series regression (see e.g. Box & Tiao, 1975), also known as a regression discontinuity in time (RDIT). There have been a number of critiques of RDIT methods (Hausman & Rapson, 2018) which we address in our

¹⁶We do exclude some period after the reform as a robustness check.

¹⁷Our estimates are not sensitive to this cleaning.

¹⁸Our estimates are not sensitive to this choice.

¹⁹Indoor-only cultivators account for approximately 50 percent of the total production market. One reassuring fact that makes us think the indoor only market is at least a plausible proxy for the entire market is that one year later outdoor production remains more than 50 percent of the market – in fact it is closer to 60 percent. If indoor growers responded to the tax reform but outdoor growers did not, we would expect a decline in the share of outdoor production. The fact that we observe the opposite suggests that, if anything, our estimates are conservative.

implementation. Most importantly, we select a relatively narrow bandwidth (measured in weeks, not years) with which to identify the effect of the reform and we provide graphical evidence to support the hypothesis that our estimated effect is indeed generated by the immediate post-reform response. Second, we aggregate our data to the weekly level to avoid the challenges in estimating day-of-week fixed effects; day-of-month effects are not important in this setting. Third, we account for the time-series nature of our data by allowing for firm-level autocorrelation between outcomes (i.e. we cluster at the firm level)²⁰ and by including lagged values of the dependent variable as a robustness check. Fourth, to ensure that compositional changes do not affect our results, we explore the balance of our panel of firms and correlated random effects (which are the equivalent of firm fixed effects for non-linear models). For our baseline analysis, we balance the panel in the neighborhood of our reform by including only firms that opened at least eight weeks before the reform and closed at least eight weeks after the reform. We consider a fully balanced panel and the inclusion of correlated random effects as robustness checks. In addition, when we examine the fraction of non-vertical activity, we always include correlated random effects because the panel is not balanced by construction – we only observe the fraction in weeks in which the manufacturer makes at least one sale to a retailer. Fifth, to provide evidence that our specifications are valid, we estimate placebo regressions that repeat our specification on year later when the cyclical trends will be similar. Combining our baseline and placebo analysis generates difference-in-RDIT estimates where we expect the second difference is approximately zero. We address other related concerns below.

²⁰Allowing for two-way autocorrelation at the manufacturer-retailer or manufacturer-week levels (Cameron et al., 2011) produces similar standard error estimates.

Our baseline regressions use the following template:

$$\log(y_{it}) = \alpha_0 + \alpha_1 TaxReform_t + \alpha_2 run_t + \alpha_3 TaxReform_t \cdot run_t + u_{it}, \quad (1)$$

where y_{it} is our outcome variable for firm i at week t , $TaxReform_t$ is a tax reform indicator that is one after July 1, 2015 and zero before, and run_t is a running variable in weeks (where zero is the week of the reform). We exclude the six weeks following the reform for the manufacturer-retailer market per the above. In our baseline analysis, we do not include an anticipation period in our regressions because the change was unexpected. Our figures provide evidence that our assumptions are plausible and we test the sensitivity of our results to these assumptions in robustness checks. Our vertical integration bandwidth is 16 weeks and the production bandwidth is 32 weeks (the planting data is noisier than the manufacturer data). We examine the robustness of our estimates to the selected time windows.

We estimate Equation (1) with non-linear methods that best match the characteristics of our dependent variables. One of our dependent variables – plantings – is best modeled as a linear hurdle model, which is the same as a Tobit model but allows for the covariates to have different extensive and intensive margin effects. We take the logs of plantings because they are approximately log-normally distributed (with the added benefit of allowing us to interpret the estimated coefficients on the binary regressors as semi-elasticities). Plantings sometimes take on a value of zero, so our dependent variable is the log of $1 + \text{outcome}$.²¹ We estimate our second dependent variable – the fraction of non-vertically integrated cannabis

²¹We could make the same distributional assumptions and estimate an exponential hurdle model for the outcome in levels, which would allow us to avoid the log of $1 + \text{outcome}$ transformation. The estimates are highly robust to this choice. However, we choose to present results with a logged dependent variable because they can be interpreted as semi-elasticities without having to present additional calculations in the text.

– using a fractional probit regression. We present the mean marginal effects from these specifications. We check the robustness of our results with respect to these model choices by re-estimating Equation (1) using OLS.

To convert our semi-elasticity estimates to an elasticity with respect to the intermediate good net-of-tax rate, we divide our semi-elasticity estimates by the legislated change in the net-of-tax rate in the intermediate goods market. We use the legislated price change because we are interested in how much this behavior changes with a particular change in the legislated tax rate. The legislated price changed from $P(1 - \tau)$ to P after the tax reform, which is $\tau/(1 - \tau)$ or 33.33% ($=0.25/0.75$).

We augment our baseline specification with several additional analyses. As an additional robustness check for our production estimates, we create a comparison group and use that to estimate a difference-in-RD regression. Specifically, we expect that cultivators will only increase production if they intend to sell plant material to other firms post-reform once the GRT has been eliminated. We therefore identify firms that sell plant material in the non-vertical market in the year after the reform as a “treated” group and firms which don’t as a comparison group. As firms self-select into these groups post-reform, these groups do not identify an average treatment effect, but rather an estimate of the difference in the response across those that choose to adjust their behavior and those that do not. If those in the comparison group don’t respond, and those in the treatment group do, this lends additional credence to our original identification strategy. Furthermore, the price changes driven by changes in tax rates also could potentially lead to production effects. As these effects would be relevant to all firms, this regression provides evidence of the presence or absence of these potential rate/price effects.

Equation (1) identifies the vertical integration response immediately after the GRT is removed for firms already in the market—i.e. those firms which have already decided whether to operate both as a cultivator and a manufacturer and paid the associated fixed costs. In the absence of the GRT, it may not be optimal to shut down or sell either their cultivation or manufacturing operations as the fixed costs have been paid and transaction costs may be high (e.g. if the two operations are physically integrated in the same location). Decisions of new entrants may well differ. Furthermore, there may be a substantial adjustment period for incumbent firms (e.g. they may slow new plantings to eventually enable the purchase of plant material from other firms but may have a significant number of plants already growing at the time of the reform). Thus, to examine the long-run effects of a GRT, we ask “how much ‘less vertically-integrated’ are firms which enter after the reform relative to incumbents?” We answer this question by estimating the parameters of

$$vertical_{it} = \beta_0 + \beta_1 entrydate_i + run_t + u_{it}, \quad (2)$$

where $vertical_{it}$ is the fraction of vertical cannabis transactions for manufacturer i in week t , $entrydate_i$ is an indicator for whether the firm opens before or after the tax reform, and run_t is a linear time trend. We define the entry date as the first date that the manufacturer sells to a retail firm.²² This regression is estimated on one year of data after the GRT has been eliminated. Our estimation window begins after the six week transition period immediately following the reform. This is a single difference regression, which leaves open

²²While firms begin their operations before their first sale, we don’t observe that date. This definition would tend to lead us to underestimate the long-run effect because a few of the firms that we define as entering after the reform will have already paid some of the fixed costs in a regime in which the GRT was still in place.

the possibility that the fraction of vertically integrated transactions naturally varies by firm age. In our data, firms tend to be more vertically integrated when they first open, so the variation by firm age would lead us to underestimate the true long-run effect. To address this, we estimate a second (placebo) difference: we examine only firms that entered before the tax reform and split them into two halves – those that entered in the first six months of the market opening and those that entered in the second six months of the market opening. If one were to combine these estimates, this would provide a difference-in-differences design (where we expect the second difference is approximately zero). We also include experience and week fixed effects in some specifications.

4 Results

In this section, we provide the first empirical evidence on the vertical integration and production inefficiency effects of a GRT. We pin down the effects of the GRT on non-vertically integrated activity in the short run and long-run in Sections 4.1 and 4.2 and then trace out its effects on long-run output in Section 4.3. We estimate heterogeneous effects on vertical integration and long-run production by firm size in Section 4.4.

4.1 Vertical Integration

Figure 1 illustrates how the fraction of cannabis sales to a retailer that originate from a vertically integrated cultivator-manufacturer pair changes for each firm when the GRT is eliminated. The hollow circles in the figure represent the raw average across all firms of the dependent variable for each week. The X-marks for the six weeks after the tax reform

indicate the transition period that we exclude from our baseline specification. Our plot spans 36 weeks pre- and post-reform – a substantially wider range than our actual local linear estimates – which provides additional information about the global behavior of our outcomes. For this reason, the solid line is a local polynomial plot of the raw data (leaving out the six week transition window), rather than a plot of our local linear regression line.

Before the tax reform, cultivators and manufacturers had a strong incentive to engage in vertically integrated activity due to the 25% GRT that was avoidable if the cultivator to manufacturer transfer were within the same firm. Empirically, we see in Figure 1 this incentive was very strong: about 94% of all activity was vertically integrated pre-reform. Table 3 presents our baseline estimates of Equation (1). We estimate in Column (1) that the elimination of the GRT leads to a 4.6 percentage point decrease in the fraction of vertically integrated sales (with a p-value of 0.018). This is a 4.90 percent decrease in vertically integrated activity in response to a 33.3 percent increase in the intermediate good net-of-tax rate, implying a short-run elasticity of vertical integration of -0.15.²³

RDIT specifications are often subject to criticism based on the possibility that estimated effects are due to cyclical or mis-specification error. Figure 1 provides compelling evidence that our results are not due to a brief cyclical effect. Moreover, the second row of Table 3 reports the result of a placebo analysis in which we conduct our estimation using data from one year after the reform date. As expected if the regression is well-specified in this context, we find the estimate is approximately zero.

We consider additional robustness checks in the remaining columns of Table 3 and in

²³The 4.90 percent decrease is calculated as: $.046/0.938$ where $.046$ is the estimated decrease in vertically integrated activity from Column (1) and 0.938 is the fraction of vertically integrated transactions pre-reform.

Figure 2. Column (2) reports OLS coefficients instead of mean marginal effects from a fractional probit model. Column (3) drops the correlated random effects. Column (4) allows for a two-week anticipation period before the tax reform by dropping the two weeks before the tax reform as well. Column (5) adds another three weeks to the transition period after the tax reform. None of these changes have a substantial effect on the estimates. In Column (6), we restrict the sample so our panel is fully balanced. This also has no meaningful effect on our estimates. Figure 2 plots the coefficient estimates for Table 3 Column (1) across many bandwidths. The estimates appear stable across the wide range of bandwidths considered.

In Table 3 Column (7), we consider the possibility that tax evasion (or other reporting inaccuracies) biases our estimates. Our estimates would be biased if firms engaged in these activities for their non-vertically integrated transactions more frequently before the tax reform when the taxes on non-vertically integrated transfers were non-zero. We do not observe tax evasion directly, but we do have data on frequent non-random audit visits and any violations firms received as a result of these visits. We drop any firm that experiences a reporting violation during the six months before or after the tax reform. The estimates and standard errors remain approximately the same.²⁴

The last column of Table 3 includes two lags of our dependent variable to take into account the autoregressive structure of the fraction of vertical transactions.²⁵ This allows us to more accurately capture the ongoing adjustment of firms post-reform given the autoregressive nature of the data. Specifically, we find that immediately after the reform, the the fraction of vertically integrated transactions decrease by 3 percentage points. Eventually, in steady

²⁴We have also confirmed that there are no changes in the audit visit and violation rates as a result of the tax reform. These results are available from the authors upon request.

²⁵Additional lags are marginally or not significant and do not change the findings discussed here.

state, the decline is 7.4 percentage points.²⁶

4.2 Vertical Integration in the Long Run

Up to this point, our analysis has focused on vertical integration behavior in the immediate aftermath of the GRT elimination. To determine the ultimate consequences of a GRT, we want to ascertain what happens in the long-run. While we can observe in Figure 1 that the fraction of vertical weight continues to decrease after the reform (and was fairly flat in the lead up to the reform), it is exceedingly difficult to know whether this is due to the autoregressive nature of the data, a longer-run response to the elimination of the GRT when firms are able to overcome fixed costs, or other secular factors. To examine this, we report estimates of Equation (2).

The top panel of Figure 3 illustrates a decline in vertically integrated activity by firms that were already in the market (marked by gray hollow circles), which we estimated in Section 4.1. This decline mostly reflects manufacturers adjusting the sourcing of their cannabis plant material to include more cultivated by other firms. We estimate in Table 4 Column (1) that vertically integrated activity among new firms that entered the market after the reform (marked by diamonds in Figure 3) is an additional five percentage points lower.²⁷ The gap in the fraction of vertically integrated sales between firms that entered before and after the reform is stable through the entire year for which we plot the data.²⁸ The persistence of this gap suggests that existing firms at the time of the reform are not slowly transitioning towards

²⁶ $0.074=0.03/(1-0.386-0.208)$.

²⁷These statements are about what happens after the six week transition period; in the case of the new entrants, note that they tend to have more vertically integrated cannabis for the first four to six weeks before reaching a stable production process.

²⁸Although not plotted, this relationship holds in approximately the same fashion throughout the data.

a mix of vertically and non-vertically integrated cannabis that looks more like new firms, and, instead, the fixed costs existing firms paid upon entry continue to largely determine the vertically integration decisions of these existing firms. Thus, the market may never ‘recover’ from the policymakers initial choice of a GRT in the sense that this market will likely always be more vertically integrated than it would have been in the absence of the original choice of a GRT. It’s also worth noting that even amongst new firms, more than 85% of transactions are vertically integrated, which suggests that it is optimal to have fairly high levels of vertical integration even in the absence of a GRT; this will dampen the production inefficiency consequences in this market relative to other markets where the optimal level of vertical integration is lower.

Table 4 Column (2) adds experience fixed effects so that we are identifying the gap from variation among same-aged firms. We also replace the time trend with weekly fixed effects to control more flexibly for time. As expected, the point estimates are slightly larger in magnitude but are also somewhat less-precisely estimated because we have absorbed some of our identifying variation. We find that post-reform entrants are 7.5 percentage points less vertically integrated than those that entered before. The long-run response can then be calculated as the sum of the short-run response for firms that were already in the market, as estimated in Table 3 Column (1), plus the additional response of new firms that entered after the GRT was eliminated. After the elimination of the GRT, firms are 12.1 percentage points less vertically integrated in the long-run ($=.046+.075$). This is a 12.9 percent decrease in vertically-integrated activity in response to a 33.3% increase in the intermediate good net-of-tax rate, implying a long-run elasticity of vertical integration of -0.39.²⁹

²⁹The 12.9 percent decrease is calculated as: $0.121/0.938$ where 0.121 is the estimated decrease in vertically

Table 4 Columns (3) and (4) and the bottom panel of Figure 3 consider another outcome – how does the likelihood a firm chooses to be completely dis-integrated in response to the eliminating of the GRT; that is, how does the probability that a manufacturer chooses not to also be a cultivator change in response to the reform. The bottom panel of Figure 3 plots the probability that a firm has no vertically integrated activity over time. The gray circles mark firms that opened before the GRT was eliminated and the green diamonds mark post-reform entrants. We interpret this plot as providing evidence in favor of two assertions: (1) there is no clear effect of the elimination of the GRT on existing firms, as expected, and (2) new firms are more than twice as likely to be completely de-integrated (i.e. operate as a manufacturer purchasing all cultivated cannabis from other firms) as existing firms. In Column (4), we estimate that those entering after the tax reform are 8.7 percentage points less likely to sell at least some cannabis that was cultivated by the same firm in a given week. This is from a base of 96.5 percentage points, implying a 9 percent decrease in the probability of operating as both a manufacturer and cultivator. This implies an elasticity of 0.27.

In the second row of Table 4, we provide a placebo test for each of our estimates – they are all near zero providing additional evidence that our specification is valid.

4.3 Production

In this section, we consider whether the shift away from vertically integrated activity after the elimination of the GRT led to increased production. To examine the immediate response, we examine at the first stage of production – plantings – which would be the first to adjust in response to the tax reform. Any change we measure in plants will then propagate through integrated activity from Column (2) and 0.938 is the fraction of vertically integrated transactions pre-reform.

to the rest of the market in the form of intermediate and final goods.

Table 5 reports estimates of Equation (1) when plantings are the outcome variable. Our baseline specification is reported in Column (1). The relevant data is plotted in the top panel of Figure 4. We find that the elimination of the GRT leads to a 22.9 percent increase in plantings or, in other words, a 22.9 percent increase in production. This is in response to a 33.3 percent increase in the net-of-tax rate, which implies a production elasticity of 0.68.

One may worry that our planting estimates are driven by cyclical or secular trends. These concerns are partially addressed by the placebo estimates provided underneath the main estimates in Table 5, which are approximately zero. Additionally, Figure 4 indicates a clear, permanent shift in plantings immediately following the reform.

Despite this evidence, it is possible that other contemporaneous events or factors influenced planting behavior to generate this result. We address this concern by isolating firms that sell in the non-vertical market post-reform—the group that benefits most from the removal of the GRT and therefore the group for which the response should be strongest—from other firms. The bottom panel of Figure 4 illustrates the data split by this grouping. Before the tax reform, cultivators who participate in the non-vertical market after the reform (gray hollow circles and black solid line) planted less than those that do not (tan triangles and black dashed line). Only these participating firms responded to the elimination of the GRT: the production gap is erased immediately upon the elimination of the GRT and over time the production in this group is higher than in the group of firms that do not participate. This suggests that those integrated firms that participate on the supply side of the non-vertical market after the elimination of the GRT have less manufacturer capacity than production capacity. The cost disadvantage created by the GRT prevented these firms from selling their

plant material to other firms and therefore these firms produced final goods at an inefficiently low level. This graphical evidence is confirmed by the estimates reported in Column (2) of Table 5. Firms that participate in the non-vertical market increased their production by 30.3 percent more than other firms (statistically significant at the 1 percent level), and firms that do not participate increased their production by a small, statistically insignificant amount.

Another potential concern regarding our baseline production estimates is that cultivators and manufacturers are responding to higher after-tax prices on the sales from manufacturers to retailers post-reform (Hansen et al., 2020), rather than responding to the change in the tax structure away from a GRT. This incentive should affect all cultivators equally; the fact that Table 5 Column (2) finds production significantly expanding only for firms that participate in the non-vertical market post-reform provides evidence that this is not the primary mechanism behind the results in this section.

In addition to certain firms producing intermediate goods at an inefficiently low rate, it is possible that other firms might choose to produce intermediate goods at an inefficiently high rate. If so, those firms that purchase plant material on the non-vertically-integrated market after the elimination of the GRT should decrease their own plant production. We estimate the effect separately for this group in Column (3) of Table 5. The estimated coefficient on the interaction between the indicator for purchasing in the non-vertical market and the post-reform indicator is small and imprecisely estimated, suggesting that this is a rare phenomenon at best. Typically, firms specialize in the non-vertical market either as sellers or buyers (see Appendix A); hence, when we combine the interaction terms from Columns (2) and (3) in Table 5 Column (4), it is not surprising that the estimates on the interaction terms are similar.

Table 6 and Figure 5 provide additional robustness checks for our baseline estimate, which are analogous to those we conducted for our analysis of the manufacturer market in Section 4.1. Column (1) repeats the estimate from Column (1) in Table 5. Column (2) estimates an OLS regression model instead of a hurdle model. Column (3) includes correlated random effects. Column (4) excludes two weeks before the tax reform to allow for an anticipation response. Column (5) excludes six weeks after the tax reform to allow for a delayed response. Column (6) restricts the regression to firms that were in business across all 64 weeks included in the regression so that the panel is fully balanced. Column (7) excludes all firms that receive a traceability violation at any point during our analysis window.³⁰ Column (8) adds two lags of the dependent variable. The estimates are highly robust to these different specifications. In Figure 5, we consider a wide range of different bandwidths and confirm that our estimates are robust to this choice.

4.4 Heterogeneity by Firm Size

Washington’s cannabis market, as with many markets, is characterized by a heterogeneous distribution of firm sizes. Indeed, if firms on the supply side of the non-vertical market for plant material are characterized by greater cultivation capacity than production capacity, it is possible that firms on the demand side of the market for plant material may have greater production capacity than cultivation capacity. In this case, we should expect the removal of the GRT to disproportionately affect firms with lower cultivation capacity. In this subsection, we consider the differences in the responses to the removal of the GRT according to the size

³⁰As above, we find similar estimates if we drop firms with any violation and we have confirmed that there are no significant changes in the audit or violation rate for cultivators after the tax reform.

of the firm. We categorize firms according to the tier of their cultivation license, which, as described in Section 2.2, effectively identifies a capacity constraint and, as seen in Table 2, is correlated with both cultivation and manufacturing activity. Measuring firm size by tier has several advantages over other measures of firm size such as average production or revenue over a given period of time: tiers capture permanent capacity constraints rather than current and/or cyclical endogenous production choices. Furthermore, tiers provide a measure of firm size for entrants before production begins.

Figure 6 plots point estimates and confidence intervals when we re-estimate our baseline result in Table 3 Column (1) for firms in each tier separately. For this analysis, we combine Tier 0 and Tier 1 firms due to the small number of Tier 0 firms (firms without a cultivator). As expected, the largest increase in non-vertically integrated weight occurs among the smallest manufacturers (Tier 0 and Tier 1). For these firms, the fraction of non-vertical activity increases by more than 15 percentage points. The effect declines substantially for the larger tiers; the estimates are slightly smaller than our baseline estimate for all firms and are not significantly different from zero. Recall that overall we do not expect a substantial increase in total weight sold 6 weeks after the tax reform because firms have to grow more cannabis plants first; however for these small firms, who could now increase their processing activity by purchasing from other cultivators, the increase in non-vertical activity led to approximately an 80 percent increase in overall activity (i.e. sales from manufacturers to retailers).

Finally, we examine the heterogeneity in the production response in Figure 7 by re-estimating our baseline planting estimates from Table 5 Column (1) separately by cultivator tier. We find that, on average, cultivators of all firm sizes increase their production by a

similar percent.³¹ This means that across all firm sizes there are some cultivators that were manufacturer-capacity-constrained (either all the time or at peak growing times). These firms increase production once that they have the opportunity to sell to other firms.

5 Conclusion

Given the increased interest in gross receipts taxes by many jurisdictions, an empirical assessment of the potential negative effects of the GRT, or any cascading tax, on vertical integration and production inefficiencies is long overdue. This analysis has previously been difficult to perform due to the infeasibility of obtaining detailed data on the entire vertical supply chain for an industry or industries affected by a GRT.

We provide the first assessment of the vertical integration and production effects of a GRT by analyzing a unique natural experiment in the Washington state cannabis industry, where the tax regime was changed from a GRT to a single excise tax at retail. We use novel data on the universe of cannabis production and sales within the state to trace out the entire vertical supply chain and understand how market participants respond to the elimination of the GRT.

We find that the GRT increased vertical integration and, in turn, decreases production. We characterize a very generalizable mechanism: the production of a vertically integrated firm is constrained by its smallest link. For example, suppose a vertically integrated automobile firm manufactures both tires and cars and its tire production is at capacity. If the firm wishes to expand car production, the firm may acquire tires from tire manufacturers or

³¹The Tier 3 estimates are somewhat imprecisely estimated reflecting the fact that a lot of Tier 3 firms are outdoor firms and are thus not in our indoor-only sample.

pay the fixed costs to increase its tire manufacturing capacity. In the presence of a GRT, the firm effectively faces a higher cost of purchasing intermediate goods and so therefore may choose not to expand car production even when capacity is available. Furthermore, as the fixed costs associated with increasing tire manufacturing require available cash-on-hand or available financing, these negative effects of a GRT may be larger for smaller and minority-owned firms (Bates et al., 2018).

Indeed, we find that the existence of a GRT hurts smaller firms with no or limited intermediate production capacity the most – manufacturers tied to small cultivators increase their production more after the removal of the GRT than manufacturers tied to larger cultivators. Under a broad GRT regime, such as those imposed by the tax policies of several states, these effects could manifest in numerous ways. Larger firms can afford to bring additional tasks in-house to avoid the higher prices generated by the additional layers of taxes created by a GRT. For example both small firms and large firms require legal services, but a small firm may not have enough work to justify hiring a lawyer full-time and must therefore pay a higher price for legal services under a GRT than a larger firm with a greater demand for similar legal services. Policymakers considering the imposition of a GRT should consider the distribution consequences, particularly given the general policy focus on supporting small businesses (see e.g. Audretsch, 2003).³²

Our estimates are likely lower bounds relative to what we would find for a GRT on all

³²Our findings are relevant not only for policymakers considering a GRT generally, but also particularly for the regulatory structure of the cannabis market as more states legalize, given the diverse regulations currently present across the United States. Colorado requires vertical integration (retail firms must produce at least 70% of the cannabis they sell) and Washington allows some types of vertical integration (between cultivators and manufacturers) while forbidding integration between retailers with firms involved in cannabis production. Our findings suggest that requiring vertical integration alters efficiency – the direction depends on the sign of externalities, which remains uncertain for cannabis (Hansen et al., 2020; Mace et al., 2020) – and hurts the smallest firms.

activity in a given jurisdiction for three main reasons. First, our empirical setting is a fairly vertically integrated industry even without a GRT, so the costs of implementing a GRT are lower. Second, because this is an industry-specific tax, investment goods like machines and equipment used in the production of cannabis are not taxed. Third, our setting features only one intermediate good sector that is taxed and in which firms can choose to vertically integrate (cultivators). In other settings with additional stages of production that may choose whether to vertically integrate, the effect of a GRT will be larger. Furthermore, firms engaged in interjurisdiction competition (such as firms which compete in integrated interstate markets) may be incentivized to simply relocate production in response to a GRT.

References

- Acemoglu, D., Griffith, R., Aghion, P., & Zilibotti, F. (2010). Vertical integration and technology: Theory and evidence. *Journal of the European Economic Association*, 8(5), 989–1033.
- Acemoglu, D., Johnson, S., & Mitton, T. (2009). Determinants of vertical integration: Financial development and contracting costs. *The Journal of Finance*, 64(3), 1251–1290.
- Agrawal, D. R. & Zimmerman, L. (2019). Production and evasion responses with limited state capacity: Evidence from major tax reforms in India. Working Paper S-89411-INC-1, International Growth Centre.
- Aizpurua-Olaizola, O., Soydaner, U., Öztürk, E., Schibano, D., Simsir, Y., Navarro, P., Etxebarria, N., & Usobiaga, A. (2016). Evolution of the cannabinoid and terpene content during the growth of cannabis sativa plants from different chemotypes. *Journal of Natural Products*, 79(2), 324–331. PMID: 26836472.
- Alfaro, L., Conconi, P., Fadinger, H., & Newman, A. F. (2016). Do Prices Determine Vertical Integration? *The Review of Economic Studies*, 83(3), 855–888.
- Allen, D. W. & Lueck, D. (2019). The organization of vineyards and wineries. In *The Palgrave handbook of wine industry economics* (pp. 325–337). Springer.
- Amin, M. R. & Ali, D. W. (2019). *Pharmacology of Medical Cannabis*, (pp. 151–165). Cham: Springer International Publishing.
- Atalay, E., Hortasu, A., & Syverson, C. (2014). Vertical integration and input flows. *American Economic Review*, 104(4), 1120–48.
- Audretsch, D. B. (2003). Standing on the shoulders of midgets: The US Small Business Innovation Research program (SBIR). *Small Business Economics*, 20(2), 129–135.
- Bates, T., Bradford, W. D., & Jackson, W. E. (2018). Are minority-owned businesses underserved by financial markets? Evidence from the private-equity industry. *Small Business Economics*, 50(3), 445–461.
- Best, M. C., Brockmeyer, A., Kleven, H. J., Spinnewijn, J., & Waseem, M. (2015). Production versus revenue efficiency with limited tax capacity: Theory and evidence from Pakistan. *Journal of Political Economy*, 123(6), 1311–1355.
- Blair, R. D. & Kaserman, D. L. (2014). *Law and economics of vertical integration and control*. Academic Press.
- Box, G. E. P. & Tiao, G. C. (1975). Intervention analysis with applications to economic and environmental problems. *Journal of the American Statistical Association*, 70(349), 70–79.
- Bresnahan, T. & Levin, J. (2012). *Vertical Integration and Market Structure*, (pp. 853–890). Princeton University Press.

- Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2011). Robust inference with multiway clustering. *Journal of Business & Economic Statistics*, 29(2), 238–249.
- Carrillo, P., Pomeranz, D., & Singhal, M. (2017). Dodging the taxman: Firm misreporting and limits to tax enforcement. *American Economic Journal: Applied Economics*, 9(2), 144–64.
- Caulkins, J. P. (2010). Estimated cost of production for legalized cannabis. Technical report, RAND Drug Policy Research Center.
- Coase, R. H. (1937). The nature of the firm. *Economica*, 4(16), 386–405.
- Cole, J. M. (2013). Memorandum for all United States attorneys. <https://www.justice.gov/iso/opa/resources/3052013829132756857467.pdf>. Accessed: 2017-07-20.
- Corsinovi, P. & Gaeta, D. (2019). Introduction: Regulations in the wine sector. In *The Palgrave Handbook of Wine Industry Economics* (pp. 249–252). Springer.
- Crawford, G. S., Lee, R. S., Whinston, M. D., & Yurukoglu, A. (2018). The welfare effects of vertical integration in multichannel television markets. *Econometrica*, 86(3), 891–954.
- Diamond, P. A. & Mirrlees, J. A. (1971). Optimal taxation and public production I: Production efficiency. *The American Economic Review*, 61(1), 8–27.
- Ernst & Young (2020). Worldwide corporate tax guide 2020. Technical report, Ernst & Young.
- Fan, Y. & Yang, C. (2020). Competition, product proliferation, and welfare: A study of the US smartphone market. *American Economic Journal: Microeconomics*, 12(2), 99–134.
- Green, G., Kryptonite, S., Chimera, B., & Ralpheme, R. (2001). *The cannabis grow bible*. Greg Green.
- Hansen, B., Houghton, K., Miller, K., & Weber, C. (2020). Getting into the weeds of tax invariance. *National Bureau of Economic Research Working Paper No. 23632*.
- Hansen, B., Miller, K., Seo, B., & Weber, C. (2020). Taxing the potency of sin goods: Evidence from recreational cannabis and liquor markets. *National Tax Journal*.
- Hansman, C., Hjort, J., León-Ciliotta, G., & Teachout, M. (2020). Vertical integration, supplier behavior, and quality upgrading among exporters. *Journal of Political Economy*, 128(9), 3570–3625.
- Hausman, C. & Rapson, D. S. (2018). Regression discontinuity in time: Considerations for empirical applications. *Annual Review of Resource Economics*, 10(1), 533–552.
- Hortaçsu, A. & Syverson, C. (2007). Cementing relationships: Vertical integration, foreclosure, productivity, and prices. *Journal of Political Economy*, 115(2), 250–301.
- Kaeding, N. (March 28, 2017). The return of gross receipts taxes. *Tax Foundation*.

- Kikuchi, T., Nishimura, K., & Stachurski, J. (2018). Span of control, transaction costs, and the structure of production chains. *Theoretical Economics*, 13(2), 729–760.
- Kopczuk, W. & Slemrod, J. (2006). Putting firms into optimal tax theory. *American Economic Review*, 96(2), 130–134.
- LaCorte, R. (June 30, 2015). Gov. Inslee signs recreational marijuana reform bill. *Associated Press*.
- Mace, C., Patel, E., & Seegert, N. (2020). Marijuana taxation and imperfect competition. *National Tax Journal*, 73(2), 545–592.
- McClure, C. (March 28, 2017). Why Ohio should not introduce a gross receipts tax – Testimony on the proposed Commercial Activity Tax. *State Tax Notes*.
- McGowan, D. (2017). Digging deep to compete: Vertical integration, product market competition and prices. *The Journal of Industrial Economics*, 65(4), 683–718.
- Mikesell, J. L. (2007a). Gross receipts taxes in state government finances: A review of their history and performance. *Tax Foundation Background Paper No. 53*.
- Mikesell, J. L. (2007b). State gross receipts taxes and the fundamental principles of tax policy. *State Tax Notes*.
- Miller, K. & Seo, B. (2018). Tax revenues when substances substitute: Marijuana, alcohol, and tobacco. *Kelley School of Business Research Paper*, (18-27).
- Phillips, A. & Ibaid, M. (2019). The impact of imposing sales taxes on business inputs. Technical report, Ernst & Young LLP.
- Pogue, T. (2007). The gross receipts tax: A new approach to business taxation. *National Tax Journal*, 60(4), 799–819.
- Ring, Jr., R. J. (1999). Consumers share and producers share of the general sales tax. *National Tax Journal*, 52(1), 79–90.
- Salinger, M. A. (1988). Vertical mergers and market foreclosure. *The Quarterly Journal of Economics*, 103(2), 345–356.
- Slemrod, J., Collins, B., Hoopes, J. L., Reck, D., & Sebastiani, M. (2017). Does credit-card information reporting improve small-business tax compliance? *Journal of Public Economics*, 149, 1 – 19.
- Smart, M. & Bird, R. M. (2009). The impact on investment of replacing a retail sales tax with a value-added tax: Evidence from Canadian experience. *National Tax Journal*, 62(4), 591–609.
- Testa, W. & Mattoon, R. (2007). Is there a role for gross receipts taxation? *National Tax Journal*, 60(4), 821–840.

Wildasin, D. E. (2001). Financing state and local government: Future challenges and opportunities. In D. E. Wildasin, M. T. Childress, M. Hackbart, L. K. Lynch, & C. W. Martie (Eds.), *Sales Taxation in Kentucky: Problems and Prospects* (pp. 27–38). Kentucky Long Term Policy Research Center.

Williamson, O. E. (1971). The vertical integration of production: Market failure considerations. *The American Economic Review*, *61*(2), 112–123.

Tables

Table 1: Gross Receipts Taxes by State

State	State Tax rate(s)	Notes
Delaware	0.0945% - 0.7468%	“Gross Receipts Tax.” 54 rates. Monthly deductions by industry: \$100K - \$1.25M.
Hawaii	0.15% - 4%	“General Excise Tax”. 2 main rates: 0.5% for wholesale, 4% for retail. + county surcharges up to 0.5%.
Kentucky	0.095%	“Limited Liability Entity Tax” C-corporations and pass-through entities pay the maximum of the GRT and a 0.75% tax on gross profit. Taxes paid are credited toward corporate income tax payments.
Ohio	0.26%	“Commercial Activity Tax”. All businesses with gross receipts above \$150,000.
Oregon	0.57%	“Corporate Activity Tax” for revenue exceeding \$1 million. The tax is \$250 + 0.57%. May deduct 35% of apportioned cost of goods sold or labor expenses. Several industries are exempt including grocery stores and utilities. Law goes into effect 1/1/2020.
Nevada	0.051% - 0.331%	“Commerce Tax” for revenue exceeding \$4,000,000 . 26 separate rates by industry.
New Mexico	5%	“Gross Receipts Tax.” + local option taxes 0.125% - 4.25%. Some “anti-pyramiding deductions” and some exemptions (e.g. grocery stores).
Texas	0.331% - 0.75%	“Franchise Tax”. 3 rates. Tax is either 70% of revenue (from federal income tax), or revenue - COGS, or revenue - wage compensation, or revenue - \$1 million.
Washington	0.13% - 3.3%	“Business & Occupation Tax”. 35 separate rates by industry.

This table only includes GRT taxes on all industries. Some states have GRT on only one industry. For example, Alabama, Florida, Kentucky, Pennsylvania, and West Virginia have GRTs on some or all utilities and/or other miscellaneous industries. Note: all states tax based on some notion of nexus (e.g. firms that are located elsewhere are still subject to the gross receipts tax for activity that occurs in that state and firms located in that state do not usually owe taxes on goods sold to other states).

Table 2: Pre-Reform Summary Statistics

Variable	Obs.	Mean	Std. Dev.
<i>Tier 1</i>			
Number of Weekly Cultivator Plantings	827	13.44	51.44
Vertical Manufacturer Weight Sold	1,159	508.98	956.23
Non-Vertical Manufacturer Weight Sold	1,159	21.66	178.65
Number of Manufacturer Transactions	1,159	5.19	9.55
<i>Tier 2</i>			
Number of Weekly Cultivator Plantings	1,394	38.44	123.71
Vertical Manufacturer Weight Sold	3,173	1,142.09	1,722.40
Non-Vertical Manufacturer Weight Sold	3,173	53.32	358.20
Number of Manufacturer Transactions	3,173	12.21	16.83
<i>Tier 3</i>			
Number of Weekly Cultivator Plantings	416	185.37	426.53
Vertical Manufacturer Weight Sold	2,512	2,518.18	4,786.74
Non-Vertical Manufacturer Weight Sold	2,512	165.77	979.45
Number of Manufacturer Transactions	2,512	25.21	43.51
An observation is a firm-week in the 16 weeks prior to the reform. It includes all firms included in the respective portions of the analysis, so for cultivators, it only includes indoor-only firms.			

Table 3: Vertical Integration Response

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Tax Reform	-0.046**	-0.045**	-0.051**	-0.051**	-0.049*	-0.040*	-0.047**	-0.030**
	(0.019)	(0.020)	(0.022)	(0.021)	(0.027)	(0.021)	(0.022)	(0.013)
Vertical _{t-1}								0.386***
								(0.030)
Vertical _{t-2}								0.208***
								(0.046)
Placebo Tax Reform	0.007	0.007	0.014	0.010	-0.001	-0.000	0.005	0.007
	(0.012)	(0.012)	(0.015)	(0.013)	(0.018)	(0.012)	(0.014)	(0.008)
Vertical _{t-1}								0.333***
								(0.033)
Vertical _{t-2}								0.219***
								(0.028)
Observations	4,271	4,271	4,271	3,947	3,788	3,503	3,279	4,271
Processor Firms	238	238	238	238	238	190	190	238
Placebo Observations	7,890	7,890	7,890	7,298	7,000	7,368	6,622	7,890
Placebo Processor Firms	437	437	437	437	437	397	366	437
Regression	Non-linear	OLS	Non-linear	Non-linear	Non-linear	Non-linear	Non-linear	OLS
Correlated Random Effects?	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes
Pre-Reform Weeks Excluded	0	0	0	2	0	0	0	0
Post-Reform Weeks Excluded	6	6	6	6	9	6	6	6
Processors Included	All	All	All	All	All	Balanced	No Viol.	All

Standard errors clustered by manufacturer organization are in parentheses. *10% significance level. **5% significance level. ***1% significance level. Each coefficient is the estimated coefficient on *TaxReform* from equation (1). The outcome is the fraction of vertical transactions for each firm-week. Other variables in equation (1) are included, but not reported. The non-linear estimates are mean marginal effects from a fractional probit regression. The top row are the estimates from the tax reform. The bottom row are placebo estimates – this analysis moves the entire regression equation one year forward in time and re-estimates the regression. All estimates in the placebo row should be approximately zero if we have a well-specified regression design.

Table 4: Long-Run Vertical Integration Response

	(1)	(2)	(3)	(4)
Enter Post-Reform	-0.050** (0.025)	-0.075* (0.045)	-0.082*** (0.021)	-0.087** (0.037)
Enter Post-Reform Placebo	0.013 (0.023)	-0.010 (0.049)	-0.004 (0.016)	-0.016 (0.033)
Observations	14,705	14,705	14,705	14,705
Processor Firms	548	548	548	548
Placebo Observations	8,954	8,954	8,954	8,954
Placebo Processor Firms	280	280	280	280
Time Trend	Yes	No	Yes	No
Week Fixed Effects	No	Yes	No	Yes
Experience Fixed Effects	No	Yes	No	Yes

Standard errors clustered by manufacturer organization are in parentheses. *10% significance level. **5% significance level. ***1% significance level. Each coefficient is the estimated coefficient on *EntryDate* from equation (2). Other variables in that equation are included, but not reported. 52 weeks are included in the regression. The estimates in Columns (1) and (2) are mean marginal effects from a fractional probit regression for the outcome fraction of non-vertically integrated transactions. The estimates in Columns (3) and (4) are mean marginal effects from a probit regression for an indicator that is one if the firm has some vertical activity that week. The top row are the estimates from the tax reform. The bottom row are placebo estimates – this analysis moves the entire regression equation 28 weeks earlier in time and re-estimates the regression. All estimates in the placebo row should be approximately zero if we have a well-specified regression design.

Table 5: Cultivator Response

	(1)	(2)	(3)	(4)
Tax Reform	0.229*** (0.075)	0.065 (0.084)	0.219** (0.088)	0.040 (0.093)
Sell Non-Vert x Tax Reform		0.303*** (0.112)		0.312*** (0.111)
Buy Non-Vert x Tax Reform			0.014 (0.116)	0.041 (0.113)
Placebo Reform	0.018 (0.055)	0.004 (0.071)	0.035 (0.059)	0.026 (0.072)
Sell Non-Vert x Placebo Reform		0.019 (0.084)		0.008 (0.082)
Buy Non-Vert x Placebo Reform			-0.056 (0.093)	-0.056 (0.091)
Tax Reform Observations	10,526	10,526	10,526	10,526
Tax Reform Producer Firms	178	178	178	178
Placebo Observations	17,115	17,115	17,115	17,115
Placebo Producer Firms	282	282	282	282

Standard errors clustered by cultivator organization are in parentheses. *10% significance level. **5% significance level. ***1% significance level. The estimated coefficient in Column (1) is the coefficient on *TaxReform* from equation (1). The outcome is the log of plantings. Other variables in that equation are included, but not reported. Columns (2) - (4) interact the tax reform variable with measures for whether the firm bought or sold in the non-vertical market in the year after the tax reform. The bandwidth is 32 weeks. The estimates are mean marginal effects from a linear hurdle model. The bottom rows are placebo estimates – this analysis moves the entire regression equation one year forward in time and re-estimates the regression. All estimates in the placebo row should be approximately zero if we have a well-specified regression.

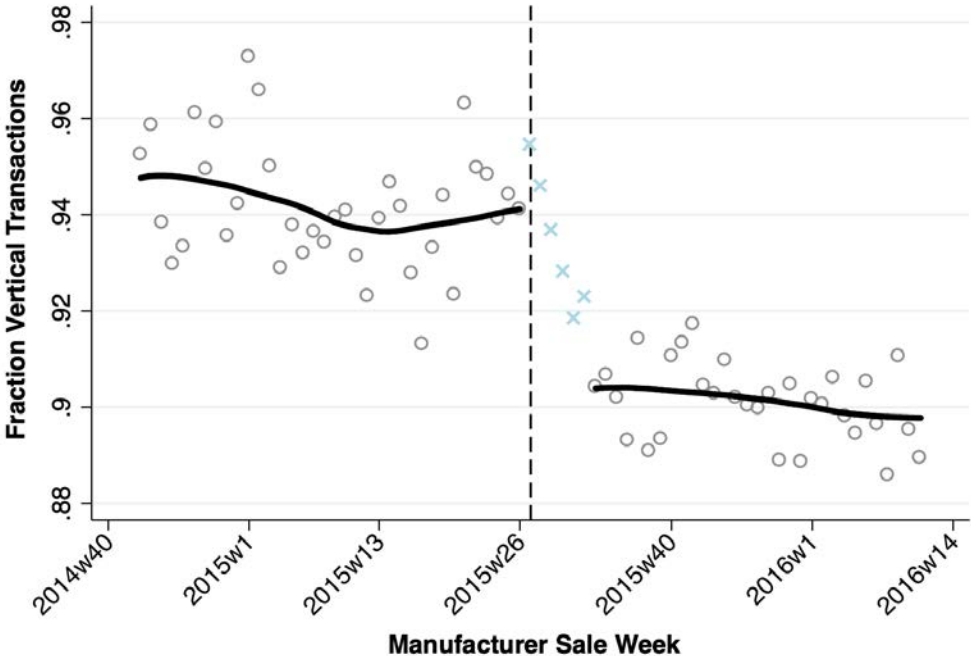
Table 6: Cultivator Response Robustness Checks

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
TaxReform	0.229*** (0.075)	0.224*** (0.074)	0.230*** (0.073)	0.202** (0.083)	0.303*** (0.090)	0.290*** (0.094)	0.192** (0.077)	0.198*** (0.069)
Vertical _{t-1}								0.017 (0.018)
Vertical _{t-2}								0.027* (0.016)
Observations	10,526	10,526	10,526	10,170	9,458	6,300	9,013	10,526
Producer Firms	178	178	178	178	178	100	153	178
Regression	Non-linear	Log OLS	Non-linear	Non-linear	Non-linear	Non-linear	Non-linear	Log OLS
Correlated Random Effects?	No	No	Yes	No	No	No	No	Yes
Pre-Reform Weeks Excluded	0	0	0	2	0	0	0	0
Post-Reform Weeks Excluded	0	0	0	0	6	0	0	0
Producers Included	8 weeks	8 weeks	8 weeks	8 weeks	8 weeks	Balanced	No Viol.	8 weeks

Standard errors clustered by cultivator organization are in parentheses. *10% significance level. **5% significance level. ***1% significance level. Each coefficient is the estimated coefficient on *TaxReform* from equation (1). The outcome is the log of plantings. Other variables in equation (1) are included, but not reported. The non-linear estimates are mean marginal effects from a linear hurdle model.

Figures

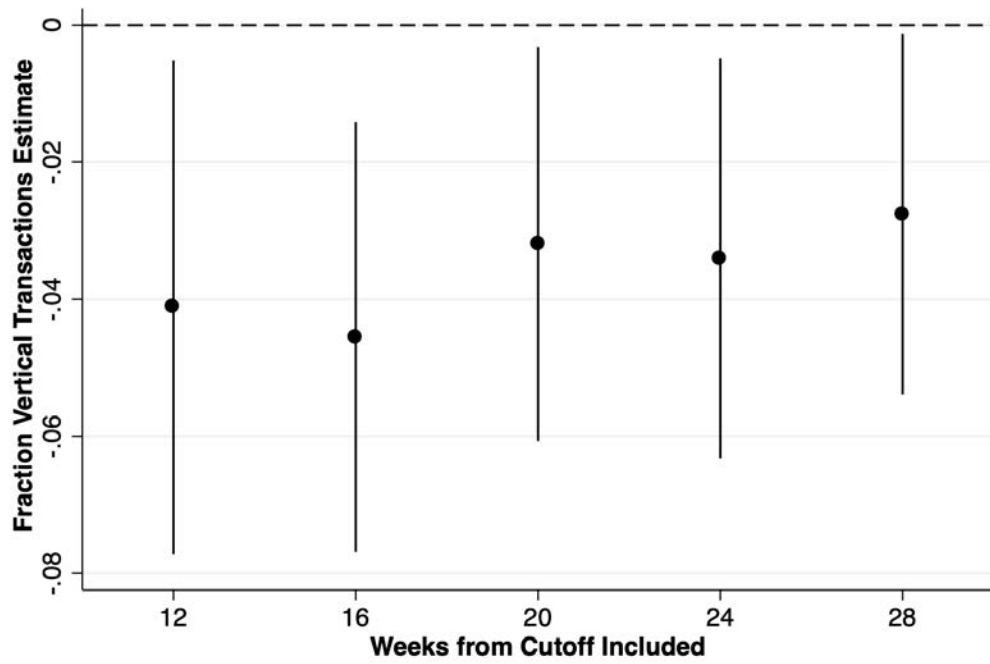
Figure 1: Vertical Integration Response



a

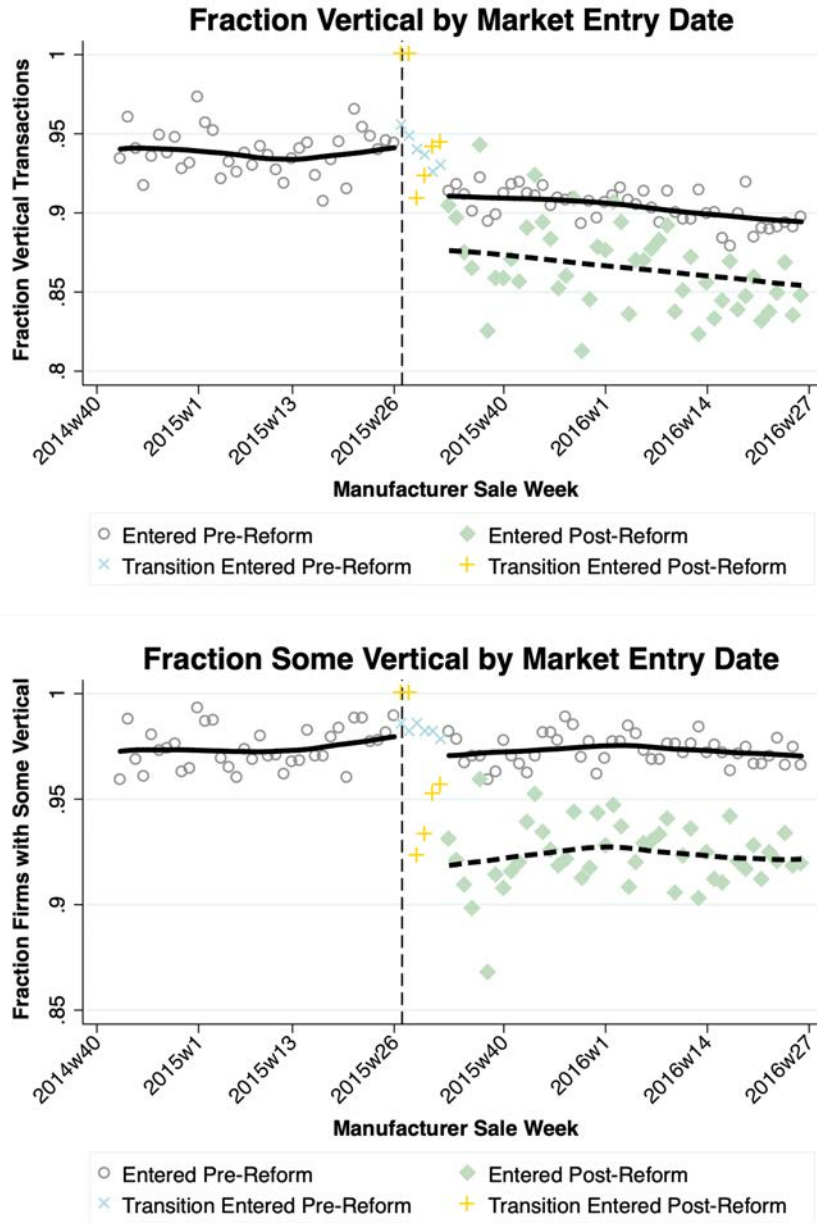
This figure plots the outcome for Table 3 Column (1). The hollow circles represent the raw average of the dependent variable by week. The blue X's for the six weeks after the tax reform indicate the transition period that we exclude from our main regressions. The solid line is a fourth-order local polynomial plot of the raw data (leaving out the six weeks transition window). The dashed vertical line marks the week of the tax reform.

Figure 2: Vertical Integration Bandwidth Choice



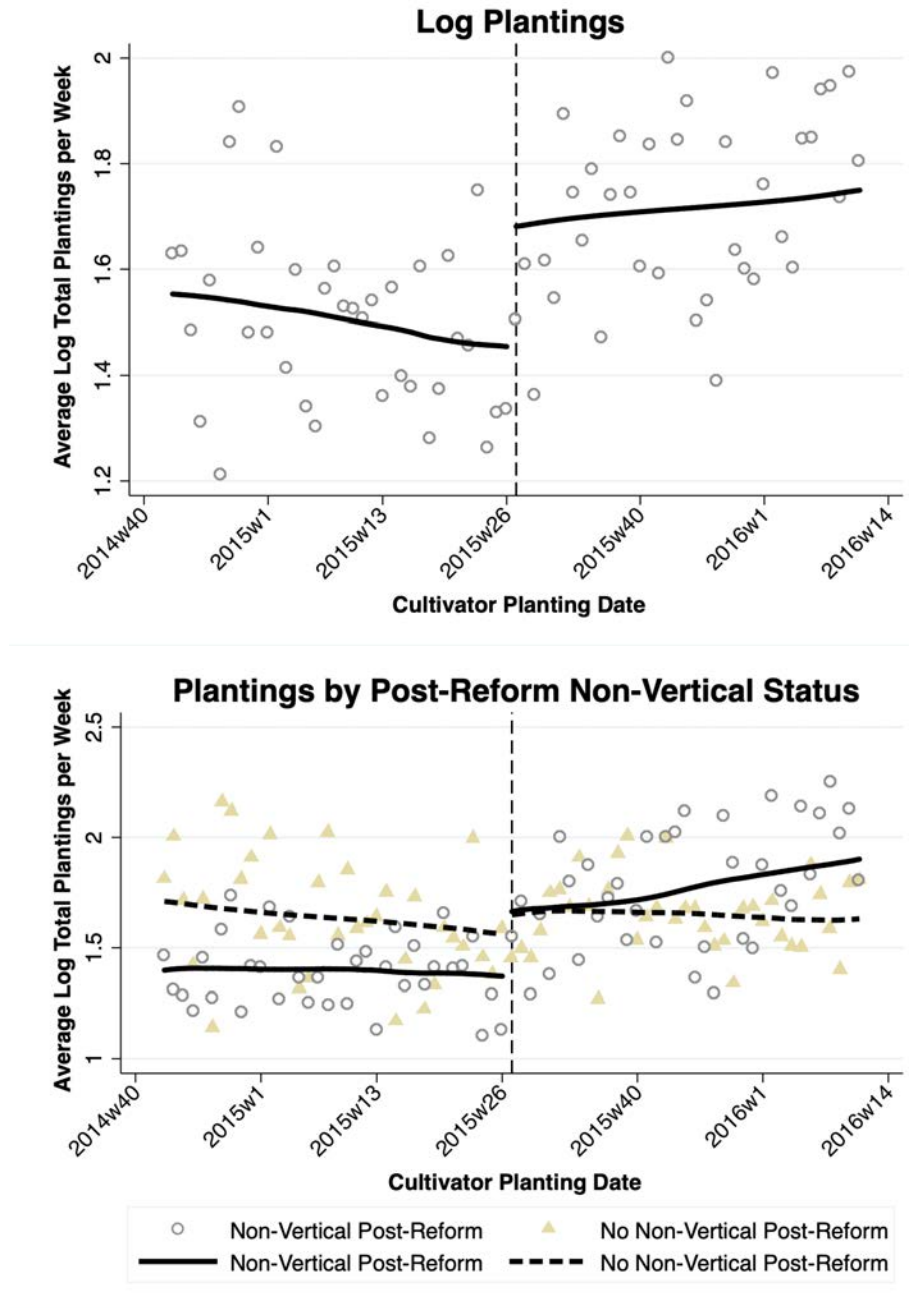
This figure plots varied bandwidths for the estimate of the coefficient on *TaxReform* in equation (1) that is found in Table 3 Column (1). Recall the main bandwidth is 16 weeks. The dots indicate the point estimates and the lines indicate 90% confidence intervals.

Figure 3: Long-Run Vertical Integration Response



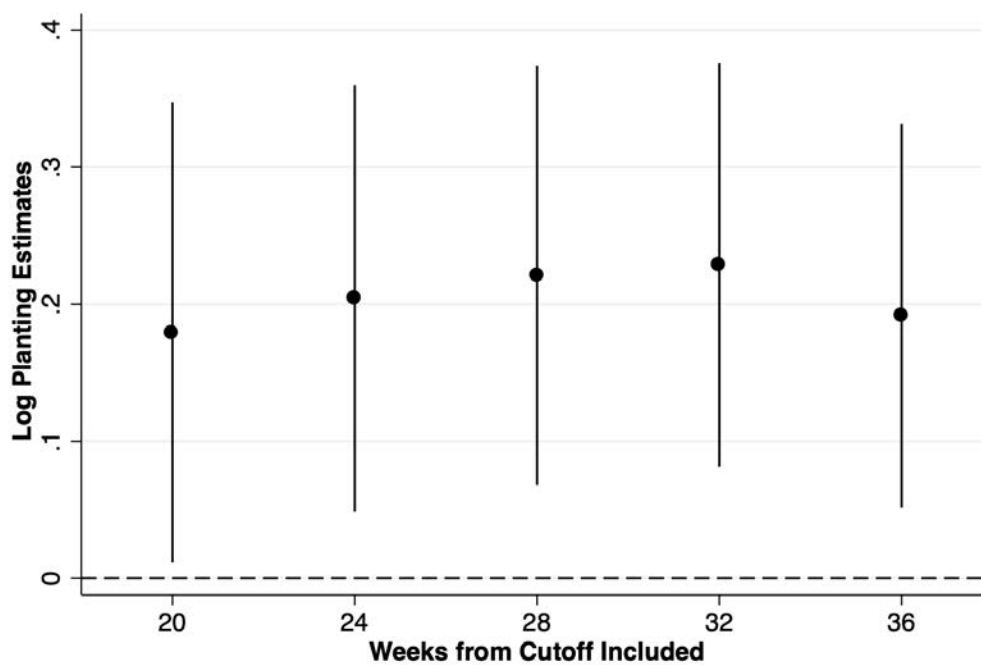
These figures plot the outcomes for Table 4 Columns (1) and (3). The hollow circles represent the raw average dependent variable by week for firms that entered before the tax reform (except the six transition weeks post-reform that are excluded from our regression analysis, which are marked by blue X's). The green diamonds represent the dependent variable by week for firms that entered after the tax reform (except the six transition weeks post-reform that are excluded from our regression analysis, which are marked by yellow +'s). The solid line is a fourth-order local polynomial plot of the raw data for firms entering pre-reform (leaving out the six weeks transition window). The dashed line is a fourth-order local polynomial plot of the raw data for firms entering post-reform (leaving out the six weeks transition window). The dashed vertical line marks the week of the tax reform.

Figure 4: Production Response




These figures plot the outcome for Table 5 Columns (1) and (2). In the top panel, the hollow circles represent the raw average weekly log plantings. The solid line is a fourth-order local polynomial plot of the raw data. The dashed vertical line marks the week of the tax reform. In the bottom panel, the hollow circles and black solid line represent the weekly log plantings for firms that participate in the non-vertical cultivator-manufacturer market in the year after the tax reform. The tan triangles and black dashed line represent the same averages for firms that do not participate in the non-vertical cultivator-manufacturer market in the year after the tax reform.

Figure 5: Production Bandwidth Choice




This figure plots varied bandwidths for the estimate of the coefficient on *TaxReform* in equation (1) that is found in Table 5 Column (1). Recall the main bandwidth is 32 weeks. The dots indicate the point estimates and the lines indicate 90% confidence intervals.

Figure 6: Vertical Integration Response by Cultivator Tier Size



This figure splits the estimate originally found in Table 3 Column (1) by cultivation tier. The dependent variable is the fraction of non-vertical transactions and the reported estimate for each tier is the coefficient on *TaxReform* in equation (1). The dots indicate the point estimates and the lines indicate 90% confidence intervals.

Figure 7: Production Response by Cultivator Tier Size



This figure splits the estimate originally found in Table 5 Column (1) by cultivation tier. The dependent variable is log plantings and the reported estimate for each tier is the coefficient on *TaxReform* in equation (1). The dots indicate the point estimates and the lines indicate 90% confidence intervals.

Appendices

A Non-Vertical Cultivator-Manufacturer Market

The vertical integration analysis in the main text focuses on the manufacturer-retail market. The main reason for this is that in that market we are able to observe *both* vertical and non-vertical transactions. However, we also have data on the non-vertical cultivator-manufacturer market – we observe sales from a cultivator to a non-vertically integrated manufacturer – and this augments our main analysis in two important ways: (1) We are able to observe shifts in firm participation in the non-vertical market immediately following the reform (that are then observable in the manufacturer-retailer market about 6 weeks later), and (2) we are able to examine whether the typical firm both buys and sells in the non-vertical market or specializes on one side of this market. Our analysis of the cultivator-manufacturer market examines flower product, which most directly maps to the “usable marijuana” category we consider in the manufacturer-retail market.

Figure A.2 plots how many cultivators and manufacturers participate each week in the non-vertical cultivator-manufacturer market. About the same number of cultivators and manufacturers participate each week. Almost no cultivator-manufacturer pairs participate in both sides of the market in the same week. The number of firms participating in the non-vertical market goes up by more than 50% in the first month after the elimination of the GRT.

Table A.1 presents statistics for the entire year before (“pre”) and after (“post”) the elimination of the GRT. Manufacturer statistics are in Panel A and cultivator statistics are in Panel B. The first row is the average number of firms that participate each week in the non-vertical cultivator-manufacturer market as plotted in Figure A.2 divided by the average number of active firms (both vertical and non-vertical) that year. Just as we saw in the figure, there is a sizeable increase after the tax reform and less than 1 percent of firms participate in both sides of the market in the same week both pre and post reform (the last

two columns of Table A.1).

The next row of Table A.1 presents the same statistics but rather than the weekly average, presents them for the percent of firms that have ever participated in the non-vertical market at some point in the entire year. Not surprisingly, these percentages are higher – there are roughly four times as many firms that participate at some point during the year relative to the weekly average. But participating in both sides of the market at some point during the year is still quite rare – only 5.9% of manufacturers do this before the reform. This translates to 18% of manufacturers who participate in the non-vertical cultivator-manufacturer market participate in both sides of the market at some point during the year. And the numbers are broadly similar post-reform. We provide the same statistics for cultivators in Table A.1 Panel B and the narrative is similar. This suggests that most firms specialize and either produce extra flower to sell on the non-vertical cultivator-manufacturer market or specialize as a manufacturer and purchase extra flower on the non-vertical cultivator-manufacturer market.

B Appendix Tables

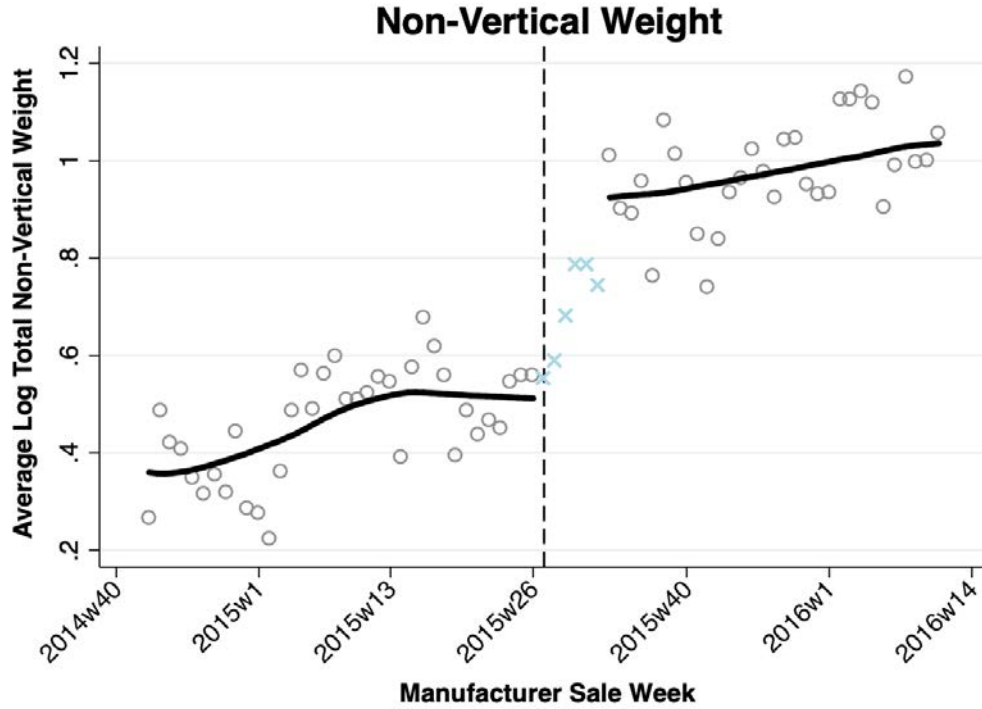
Table A.1: Cultivator-Manufacturer Market Summary Statistics

Panel A: Manufacturers in Non-Vertical Market				
	Pre	Post	Both Pre	Both Post
Weekly Average Percent of Firms:	8.95%	13.71%	0.16%	0.74%
Yearly Total Percent of Firms:	32.99%	45.19%	5.90%	6.40%
Panel B: Cultivators in Non-Vertical Market				
	Pre	Post	Both Pre	Both Post
Weekly Average Percent of Firms:	8.84%	14.23%	0.15%	0.65%
Yearly Total Percent of Firms:	38.51%	70.92%	5.50%	5.78%

These are summary statistics about participation in the non-vertical cultivator-manufacturer market. The top panel describes manufacturer activity as a percentage of the total number of firms in the market for the given time period and the bottom panel does the same for cultivators. “Pre” is 1 year before the tax change, “Post” is 1 year after the tax change. When we consider “Both” this is the percentage of firms in that week or year that participate in both sides of the market (i.e. both sell cannabis as a cultivator and buy cannabis as a manufacturer). The “Weekly Average Percent of Firms” is the average fraction of cultivators (or cultivators) that participate in the non-vertical market each week. And, “Yearly Total Percent of Firms” is the fraction of manufacturers (or cultivators) that participate in the non-vertical market at least once that year. The numerators of these two variables are from the cultivator-manufacturer data and the denominator is all manufacturers (or cultivators) that ever have a transaction sold in the manufacturer-retailer market.

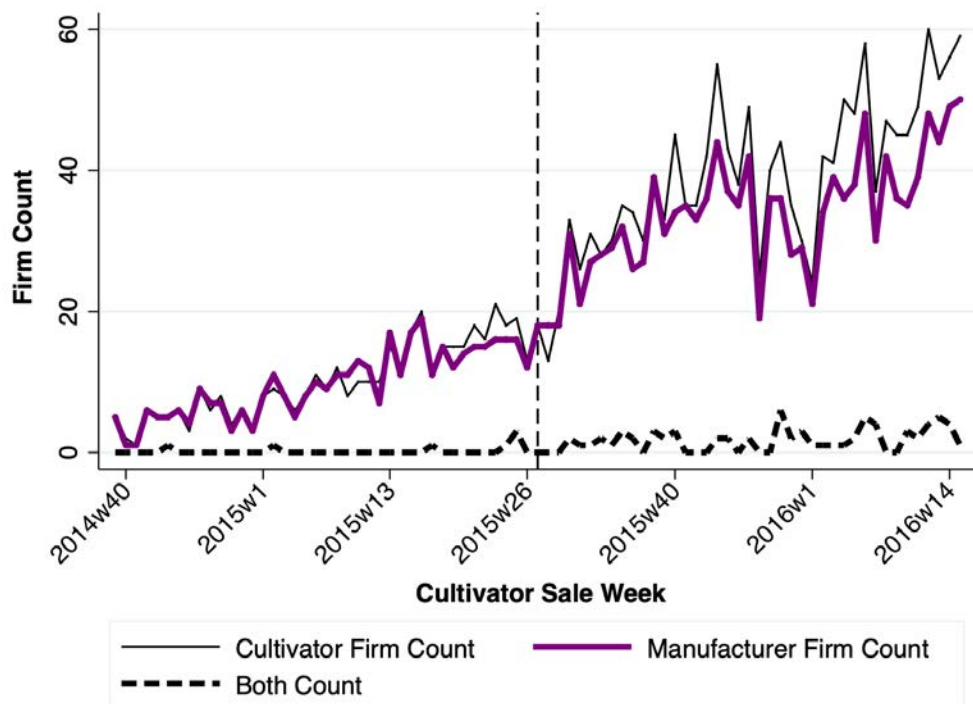
C Appendix Figures

Figure A.1: Vertical Integration Response



This figure provides an alternative dependent variable – non-vertical weight – to examine vertical integration behavior in response to the elimination of the GRT. The hollow circles represent the raw average weekly dependent variable for each week. The blue X's for the six weeks after the tax change indicate the transition period that we exclude from our main analysis. The solid line is a fourth-order local polynomial plot of the raw data (leaving out the six weeks transition window). The dashed vertical line marks the week of the tax reform. Estimating the response using equation (1), we estimate that the elimination of the GRT increases the grams of non-vertically cultivated cannabis sold by 51 percent – these are all transactions between firms that would have happened in the absence of the GRT that now do take place. Both the extensive and intensive margin are significantly contributing to the 51% increase; that is, firms are both entering the non-vertically integrated market, and conditional on entering, transacting in this market in larger quantities.

Figure A.2: Participants in the Cultivator-Manufacturer Market



This figure plots the counts of cultivators and manufacturers that participate in the non-vertical cultivator-manufacturer market each week for about 9 months (40 weeks) before and after the reform. The thin black line plots the number of cultivators each week. The thick purple line does the same for manufacturers. The thick black dashed line plots the number of firms that participate in both sides of the market in the same week. The dashed vertical line marks the week of the tax reform.