Vertical Integration and Production Inefficiency in the Presence of a Gross Receipts Tax

Benjamin Hansen*
University of Oregon, NBER, IZA

Keaton Miller*
University of Oregon

Caroline Weber*
University of Kentucky

May 2022

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 about twice as large. We find these incentives lead to large output losses – production increases by 23 percent when the GRT is eliminated.

JEL Codes: H20, H30, H70.
Keywords: Gross Receipts Tax, Turnover Tax, Vertical Integration, Natural Experiment, Cannabis, Excise Tax.

*Hansen: bchansen@uoregon.edu; Miller: keatonm@uoregon.edu; Weber: caroline.weber@uky.edu. 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."

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 that was comprised of a vertically integrated cultivator and manufacturer enjoyed a potential cost advantage over dis-integrated competitors.[2] On July 1, the GRT was

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

[2] 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.

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, 2015, and signed by the Governor on June 30.

In Section 4 we introduce a general model of vertical integration in the supply chain following the conduct parameter approach of Weyl & Fabinger (2013). We calibrate the model using pre-reform data and simulate post-reform outcomes to generate testable hypotheses.

In Section 6 we measure the effects of this reform empirically 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.[3]

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 about twice as large as the short-run effects within the incumbent group. We estimate that the long-run elasticity

[3]We 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.

of vertical integration with respect to the intermediate good net-of-tax rate is -0.25.

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 that sell intermediate goods to other manufacturers 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). We discuss the welfare implications of these results in Section 7 and conclude by discussing the policy implications of these findings in Section 8.

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.[4] 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 large number of firms in a statewide competitive market 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). Most of this literature documents policy-induced frictions that lead to *too little* vertical integration, such as financial markets (Acemoglu et al., 2009), technology differences (Acemoglu et al., 2010), or quality differences (Hansman et al., 2020). Perhaps most similar to our work, Grubert (2003), Desai et al. (2004), and Egger & Seidel (2013) document a setting in which tax evasion leads to *too much* vertical integration – corporate tax rate differentials lead to excess vertical integration among multinationals so that these firms can take advantage of international transfer pricing opportunities. We will return in Section 8 to compare the vertical integration estimates from this literature to our own estimates.

[4] See (Bresnahan & Levin, 2012) for a review.

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).[5] 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 stages of production at the same rate without allowing any deductions for costs.[6]

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).

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

[6]For 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.

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 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 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.[7] 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. However, any owner of a retail license may not have any ownership interest in any other type of license, and vice versa.[8] 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.

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 Δ9-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

[7] Washington 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.

[8] These restrictions extend to relatives of licensees.

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 analyses focus on these "usable marijuana" products exclusively, which is about 75% of the total market.[9] 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

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

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.[10] Reporting is enforced through frequent 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.[11] 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 produc-

[10] Retailers 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 significant under-reporting.

[11] Atalay 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.

tion 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 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 6.1).

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

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.[12] Once harvested, flowers and other plant material are collected and converted into a new "inventory lot" that is assigned a unique identifier (ID); products or material within a single inventory lot are assumed to be homogeneous. These intermediate products may progress through several processing steps before wholesale distribution.

The last processing step is the division of a large wholesale inventory lot 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.[13] 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.[14] Vertically dis-integrated activity originates from two sources: (1)

[12] Strains are defined by the cultivator.

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

[14] Capacity constraints or ebbs and flows in production may lead a cultivator which is part of a vertically

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.

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.[15]

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 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 firm to contract with a different manufacturer which is part of a separate vertically integrated firm.

[15]Unfortunately, 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 its 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.

integrated cannabis was similar at approximately 95%.

4 A Model of Vertical Integration

To motivate our empirical analyses of vertically integrated and disintegrated production, we introduce a stylized model of a supply chain and characterize its equilibria. We then calibrate an instantiation of the model to generate motivating hypotheses. Within the model, there are two activities needed to produce final goods: the production of intermediate goods (i.e. growing plant material) and the processing of those intermediate goods into final goods, which are then sold.[16] Firms may participate in one or both of these activities, trading in the intermediate and/or final goods market. Both markets exhibit imperfect competition characterized by a firm-specific *conduct parameter* following Weyl & Fabinger (2013) and Hansen et al. (2021) (whose exposition we follow); this structure nests many common models of imperfect competition and is appropriate here as firms in different geographies in our setting face varying levels of competition. We abstract from certain details for simplicity and to broaden the applicability of the model across industries. For example, above we argue that capacity constraints play an important role in Washington's cannabis industry; in the model we assume smooth cost functions and then consider their relation to capacity constraints.

The competitive environment consists of n firms, denoted by i. Each firm chooses a quantity of intermediate goods to produce (grow) q_{gi} and a quantity of intermediate goods to process into final goods q_{pi}—if these quantities differ, the firm must participate in the

[16] In the cannabis industry, processors sell goods to retailers who then sell goods to consumers. Our model is agnostic as to the source of demand (i.e. downstream firms vs. consumers).

intermediate goods market to be described below. Each activity $a \in g, p$ generates costs that are smooth increasing functions of the quantity of the relevant good: $c_{ai}(q_{ai})$. Define $mc_{ai}(q_{ai}) = \frac{dc_{ai}}{dq_{ai}}(q_{ai})$. For ease of notation, we write mc_{ai} without explicitly including the quantity. We therefore can write $mc_i = mc_{gi} + mc_{pi}$.

In the final goods market, demand is smooth and is given by the demand function $p_f = D(q)$ where p_f is the price of the final good and $q \equiv \sum_i q_{pi}$ is the total quantity of final goods produced. The price elasticity of demand for final goods is $\epsilon_D = (-dq/dp_f) \cdot (p_f/q)$ and is assumed to be greater than unity over the range of prices. Sellers of final goods face a tax rate τ_p, so if the unit price paid by buyers of final goods is p_f, the revenue earned by the seller is $p_f(1 - \tau_p)$. We assume that firms that participate in the final goods market choose prices in order to set their elasticity-and-tax-adjusted Lerner index equal to an exogenous firm-specific conduct parameter θ_{pi}. That is,

$$\frac{p_f - mc_i/(1 - \tau_p)}{p_f}\epsilon_D = \theta_{pi}. \tag{1}$$

Weyl & Fabinger (2013) show that when products are weak substitutes, the conduct parameter ranges from 0 (perfect competition) to 1 (monopoly).

Firms trade intermediate goods, which are assumed to be identical, at a price of p_m per unit. The revenue tax rate paid by firms selling intermediate goods is τ_m—thus the unit net-of-tax revenue of selling intermediate goods is $p_m(1 - \tau_m)$. In equilibrium, firms will either purchase intermediate goods, sell intermediate goods, or neither. Denote buyers of intermediate goods with j and sellers with k. Define $q_{bj} = q_{pj} - q_{gj}$ for buyers and

$q_{sk} = q_{gk} - q_{pk}$ for sellers.[17] Let $q_b = \sum_j q_{bj}$ and $q_s = \sum_k q_{sk}$. Market clearing requires $q_b = q_s$. As above, we model imperfect competition in this market through a conduct parameter θ_{mi} for firms selling goods in this market; firms buying intermediate goods are price takers. That is,

$$\frac{p_m - mc_{gi}/(1-\tau_m)}{p_m}\epsilon_M = \theta_{mi}, \tag{2}$$

where $\epsilon_M = -(dq_s/dp_m) \cdot (p_m/q_s)$.

4.1 Interior Equilibrium

Given conduct parameters θ and cost functions c_i, equilibrium in this market is defined by the vector (q_{ai}, p_f, p_m) which has $2i+2$ terms. In this section we derive sufficient conditions to characterize an *interior equilibrium* in which each firm i participates in both activities (i.e. $q_{ai} > 0$ for all i). We focus on interior equilibria as firms in our setting generally engage in both activities to at least some extent.

Two equilibrium conditions come from the demand equation $p_f = D(q)$ and market clearing ($\sum q_{gi} = \sum q_{pi}$). Equation (1) generates i conditions since in an interior equilibrium each firm sells final goods. The remaining conditions come from considering behavior in the intermediate goods market. Consider buyer j's cost minimization problem taking q_{pj} as given. At the margin, the (price-taking) buyer must be indifferent between purchasing more goods from the intermediate market at price p_m or producing more intermediate goods

[17] For ease of exposition, we consider firms with $q_{gi} = q_{pi}$ to be buyers of intermediate goods with $q_{bj} = 0$.

themselves at marginal cost mc_{gj}. Therefore, for each buyer j, in equilibrium,

$$p_m = mc_{gj}. \tag{3}$$

We are left with k conditions, one for each seller of intermediate goods. In equilibrium, sellers of intermediate goods must be indifferent on the margin between selling intermediate goods and processing those goods internally. Note that in equilibrium this choice does not affect the price of the final good as their intermediate good is processed either way; $dp_f/dq_{sk} = 0$. The seller's indifference condition is, therefore

$$\left[\frac{dp_m}{dq_s}q_{sk} + p_m(q_s)\right](1-\tau_m) = p_f(1-\tau_p) - mc_{pk}.$$

By solving this expression for dp_m/dq_s and applying it to Equation (2) (along with the definition of ϵ_M), we obtain k equilibrium conditions

$$p_m - \frac{mc_{gk}}{1-\tau_m} = \theta_{mk}\frac{q_s}{q_{sk}}\left[\frac{p_f(1-\tau_p) - mc_{pk}}{1-\tau_m} - p_m\right] \text{ for all } k. \tag{4}$$

This equation relates the tax-adjusted marginal profit obtained by the seller in the intermediate goods market (on the left-hand side) to the tax-adjusted marginal profit (net of opportunity costs) obtained by the seller in the in the final goods market. These profits differ according to the market power of the seller in the intermediate goods market, modulated by the fraction of intermediate goods supply provided by this seller. Note that if $\theta_{mk} = 0$ (perfect competition) this condition reduces to a price-taking condition $p_m(1-\tau_m) = mc_{gk}$ that mirrors Equation (3).

These conditions—the demand function, market clearing, and Equations (1), (3), and (4)—characterize interior equilibria. We note that depending on the nature of $D(q)$ and c_{ai}, there may be zero, one, or many such equilibria.

4.2 Calibrated hypotheses

To motivate our empirical analyses, we present a simple instantiation of the model featuring heterogeneous firms and quadratic costs (to capture the flavor of capacity constraints). We calibrate the model to moments derived from our data on Washington's cannabis industry prior to the tax reform and make predictions about post-reform outcomes.

Consider an environment with two firms, $i = 1, 2$, which face costs $c_{ai}(q_{ai}) = \beta_{ai} q_{ai}^2$. We assume that demand for final goods features constant elasticity, that is $q = k p_f^{(}\epsilon_D)$. Without loss of generality, we assume that firm 1 is the seller of intermediate goods and is a price taker in that market (i.e. $\theta_m = 0$). Firm 2 is the buyer in the intermediate goods market. We further assume that $\theta_{p1} = \theta_{p2} = \theta_p$. We therefore have seven parameters to calibrate: two demand parameters, four cost parameters, and one competitive conduct parameter.

We calibrate the demand parameters using the results of Hansen et al. (2020): we set $\epsilon_D = 1.1689$ and choose k to match the average quantity sold per UPC per week at the average wholesale price we observe in the data.[18] To calibrate the remaining five parameters, we calculate five moments from the pre-reform data: the degree of vertical integration (one minus the fraction of cannabis traded), the average price per gram charged by manufacturers to retailers, the average price per gram charged by cultivators to manufacturers, and the ratios of the quantity grown or processed by those on the sell-side of the plant material

[18] Our results are robust to this choice as our other calibration moments are quantity-level-agnostic.

market to the appropriate quantity grown or processed by those on the buy-side of the plant material market.[19] We choose parameters by minimizing the sum of squared differences between the moments in the data and the moments in the model.

The results of our calibration are reported in Table 3. Panel (a) reports the calibrated values of the parameters, while Panel (b) reports the moments in the data and in the calibrated model. Despite the parsimonious specification (particularly the simple cost function and the assumption of two firms), the model is able to closely match the moments: each model moment is within approximately 1% of the data. The cost parameters indicate that firm 1 has higher costs for both cultivating and processing cannabis than firm 2, though the gap is larger for processing than for cultivating. This is consistent with the observation in Table 2 of distinct sizes of firms according to their tiers and the corresponding hypothesis that many firms in the market are effectively capacity constrained. The calibrated conduct parameter indicates that manufacturers have little market power when selling to retailers.

After calibration, we use the model to form hypotheses about post-reform market outcomes that we then take to the data in Section 6. In particular, we remove the gross receipts tax and increase the tax on final goods to 37%, and then calculate the change in the fraction of final goods that are produced in a vertically integrated way and the (log) change in the quantity cultivated by firms participating in the intermediate goods market (i.e. $\log q_{g1}^{post} - \log q_{g1}^{pre}$ and $\log q_{g2}^{post} - \log q_{g2}^{pre}$). We report these predicted outcomes in Panel (c). The model predicts that the fraction of vertically integrated final goods will decrease by 2.2% (i.e. from 94.7% to 92.5%) and the (log) quantity cultivated by the sell-side firm

[19] When computing these ratios, we partition firms into the 'sell-side' or 'buy-side' based on their post-reform outcomes, as the initial tax regime may have discouraged firms from participating in the plant material market at all. We use pre-reform quantity choices to calculate the moments.

will increase by approximately 30.6% while the quantity cultivated by the buy-side firm will increase by 1.8%.

These predictions provide hypotheses that we can then examine empirically: we expect the removal of the GRT to reduce vertical integration and increase production, particularly for firms participating in the sell-side of the market post-reform. We note that our model abstracts from a number of real-world considerations, including (but not limited to) the degree of fixed costs involved in changing capacity constraints or setting up vertically-disintegrated supply relationships. Our point estimates therefore may not match the predicted outcomes in Table 3 precisely.

5 Empirical approach

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.[20] 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 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

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

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.[21]

We aggregate the data by firm location-week and perform a few minor data cleaning steps.[22] 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).[23] We exclusively examine the indoor-only cultivator market to eliminate cyclical effects.[24]

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 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

[21] We do exclude some period after the reform as a robustness check.

[22] Our estimates are not sensitive to this cleaning.

[23] Our estimates are not sensitive to this choice.

[24] Indoor-only cultivators account for approximately 50 percent of the total production market (although this varies some by season as can be see in Figure A.5). 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.

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)[25] 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 one year later when the cyclical trends will be similar and provide permutation tests based on estimates from many placebo weeks. Combining our baseline and placebo analysis generates difference-in-RDIT estimates where we have evidence that the second difference is approximately zero. We address other related concerns below.

Our baseline regressions use the following template:

$$y_{it} = \alpha_0 + \alpha_1 TaxReform_t + \alpha_2 run_t + \alpha_3 TaxReform_t \cdot run_t + u_{it}, \tag{5}$$

where y_{it} is our outcome variable for firm i at week t, $TaxReform_t$ is a tax reform indicator

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

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 as discussed 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 (5) 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 + outcome.[26] We estimate our second dependent variable – the fraction of non-vertically integrated cannabis – 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 (5) using OLS.

To convert our semi-elasticity estimates to an elasticity with respect to the intermediate

[26]We 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 + 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.

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. The legislated net-of-tax rate change measures the change in the legislated price due to the tax reform. 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 vertical integration behavior and those that do not. If those in the comparison group have a small response and those in the treatment group have a large response – consistent with the buy and sell side predictions from the calibrated model in Section 4 – 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 (5) 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). 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).

To examine the longer-run vertical integration effects of a GRT, we re-estimate equation (5) but we include firms that open after the reform, as well as before, and our bandwidth remains the same pre-reform, but we add an entire year of data post-reform. We then introduce a new variable, $NewFirm_i$ to denote firms that enter after the reform. We define a firm's entry date as the first date that the manufacturer sells to a retail firm.[27] The regression is specified as follows:

$$vertical_{it} = \alpha_0 + \alpha_1 TaxReform_t + \alpha_2 NewFirm_i + \alpha_3 run_t + \alpha_4 TaxReform_t \cdot run_t + u_{it}. \quad (6)$$

where $vertical_{it}$ is the fraction of vertical cannabis transactions for manufacturer i in week t, $NewFirm_i$ is an indicator for whether the firm opens before or after the tax reform, and run_t is a linear time trend. Our first parameter of interest α_1 still identifies the short-run effect for incumbent firms. Our second parameter of interest α_2 identifies how much less vertically integrated are new entrants relative to incumbent firms. If we assume new entrants would have been the same across time in the absence of a tax reform, we can attribute any decline in the vertical integration of new entrants to the reform. Under these assumptions,

[27] While 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.

we expect incumbents to look like these new entrants in the long-run. This implies that the long-run effect is the sum of the short-run effect experienced by incumbent firms (α_1) and the additional response of new entrants (α_2), which we assume all firms will experience in the long-run.

The identification of α_2 comes from a single difference, which leaves open 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 control for experience in some specifications.

6 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 6.1 and 6.2 and then trace out its effects on long-run output in Section 6.3.

6.1 Vertical Integration

Figure 1 illustrates how the fraction of cannabis sales to a retailer that originate from a vertically integrated cultivator-manufacturer pair changes 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 95% of all activity was vertically integrated pre-reform. Table 4 presents our baseline estimates of Equation (5). 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.[28] We discuss this estimate in the context of other measurements of vertical integration in Section 8.

RDIT specifications are often subject to criticism based on the possibility that estimated

[28]The 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.

effects are due to cyclicality 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 4 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. The top panel of Figure 2 provides a placebo permutation test – the figure plots placebo estimates using our baseline specification, but reassigning treatment to each week in our data, except those within six weeks of the tax reform, Black Friday, and 4/20.[29] The gray dashed line marks our estimated effect; this line is more than double the next most extreme estimate. We have 45 observations in our permutation test, so the implied p-value is 0.022 (=1/45).

We consider additional robustness checks in the remaining columns of Table 4 and the bottom panel of 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 excluding the two weeks before the reform. 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. The bottom panel of Figure 2 plots the coefficient estimates for Table 4 Column (1) across many bandwidths. The estimates appear stable across the wide range of bandwidths considered.

[29] We begin our placebo estimation in mid-Februrary 2015. Before this, the smaller sample size, driven by many firms still opening and ramping up their production, makes the estimates substantially more noisy and thus a less good comparison to the period of the reform. However, if we did include earlier periods (and periods we've excluded surrounding holidays), most estimates are still near zero, and the resulting permutation test p-value would be smaller. Hence we are reporting a more informative, but also more conservative, permutation test.

In Table 4 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.[30]

The last column of Table 4 includes two lags of our dependent variable to take into account the autoregressive structure of the fraction of vertical transactions.[31] This allows us to more accurately capture the ongoing adjustment of firms post-reform given the autoregressive nature of the data. To the extent that there is some downward adjustment in vertical integration levels during the six weeks we exclude, this specification will under-estimate the effects of the reform.[32] This specification finds that immediately after the reform the fraction of vertically integrated transactions decrease by 2.3 percentage points. Eventually, in steady state, the decline is 2.6 percentage points.[33]

To summarize, across specifications, we find evidence that the reform reduced the degree of vertical integration in the short run, and the 95 percent confidence interval of this decline

[30] We 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.

[31] Additional lags are marginally or not significant and do not change the findings discussed here.

[32] When part of the transition is excluded, the suddenly lower levels of lagged vertical integration post-reform explain part of the tax reform effect. In contrast, if this transition had been included, the first period post-reform would have vertical integration lags still at pre-reform levels and so these would not explain any of the tax-reform effect. Unfortunately, this transition cannot be included here in light of other reasons for a lagged response.

[33] 0.026=0.023/(1-0.094-0.054).

includes the decrease predicted by the model in Section 4.

6.2 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. Indeed, our model predictions can be most easily interpreted as long-run predictions, as the model does not include any transition costs. 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 (7).

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 6.1. This decline largely reflects manufacturers adjusting the sourcing of their cannabis plant material to include more cultivated by other firms (as we explore further in Figure 4 below). It then introduces an additional line post-reform – average vertical integration levels of new firms (marked by light-green diamonds). The vertical integration levels of new firms are significantly lower.[34] The gap in the fraction of vertically integrated sales between incumbent and new firms is stable through the entire year for which we plot the data.[35] The persistence of this gap suggests that incumbent firms at the time of the reform are not slowly

[34] These 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.

[35] Although not plotted, this relationship holds in approximately the same fashion throughout the data.

transitioning towards 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 incumbent firms. Thus, the market may never 'recover' from the policymakers initial choice of a GRT in the sense that this market will likely be more vertically integrated than it would have been in the absence of the original choice of a GRT for a long time. 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 5 estimates the effects depicted in Figure 3. Column (1) separately identifies the effect of the tax reform on incumbent firms and the additional effect it has on new firms (α_1 and α_2 in equation (7)). We find that incumbent firms are 4.1 percentage points less likely to be vertically integrated post-reform,[36] and new firms are 4.2 percentage points less vertically integrated than incumbent firms. The long-run response can then be calculated as the sum of these two coefficients.[37] After the elimination of the GRT, firms are 8.3 percentage points less vertically integrated in the long-run (=.041+.042; p-value = 0.011). This is an 8.9 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 -0.25.[38] Column

[36]The extensive period post-reform and our inability to include correlated random effects in this specification make this estimate less precise. These issues could also bias the estimate, but any bias appears small given the similarity of this estimate with our baseline estimates above.

[37]As discussed in Section 5, summing these two estimates presumes that the vertical integration decisions of new and incumbent firms would be similar absent the reform (and because of this, the new entrants reflect the additional long-run response for all firms once they are able to overcome their fixed costs). Our placebo analysis in the table below our main results provides evidence to support this assumption as does the stable fraction of non-vertically integrated firms in the period prior to the reform shown in Figure A.3.

[38]The 8.9 percent decrease is calculated as: 0.083/0.938 where 0.083 is the estimated long-run decrease in

(2) adds controls for the log of experience and time to closing to the regression specification. As expected, the point estimates are almost identical but are also somewhat less-precisely estimated because we have absorbed some of our identifying variation. In the second row of Table 5, we provide a placebo test for each of our estimates – they are all near zero providing additional evidence that our specification is valid.

To understand some of the underlying decisions driving the estimates in Figure 3 and Table 5, we consider Figures 4, A.3, and A.4. The top panel of Figure 4 provides a histogram of the fraction of vertically integrated sales by firm for three months before (shaded gray) and after the reform (hollow green), for all incumbent firms. We see from this figure that, in response to the reform, more than 70% of existing firms didn't deviate from being mostly (or exclusively) vertically integrated. However, about 10% of firms switched from being mostly vertically integrated towards a substantial mix of vertically integrated and dis-integrated activity (i.e. they started trading more with other firms). There was no increase in fully dis-integrated firms among incumbent firms. The bottom panel of Figure 4 replicates the histogram from the top panel for incumbent firms in the post-reform period (shaded gray histogram) and adds a histogram for the same time period for all new firms (hollow green histogram). We see new firms making more bi-modal choices than incumbent firms; new firms choose to be mostly vertically integrated at about the same rate as incumbent firms, but they choose to be mostly vertically dis-integrated at a much higher rate than incumbent firms. And new firms rarely choose to mix vertically integrated and dis-integrated activities. By assuming that the new firms are a good proxy for existing firm decisions in the long-run,

vertically integrated activity from Column (1) and 0.938 is the fraction of vertically integrated transactions pre-reform.

we are effectively presuming that existing firms that currently offer a mix of vertically and non-vertically integrated cannabis will eventually transition to being fully dis-integrated. To the extent that they do not, this market will likely not reach the vertical dis-integration levels that we have documented for new firms.

Figures A.3 and A.4 provide evidence of quasi-permanent manufacturer and cultivator decisions regarding whether to be vertically integrated over time. The bottom panel of Figure A.3 shows that the fraction of stand-alone manufacturers (i.e. those without an operating cultivator) is quite stable in the pre-reform period and then increases substantially in the period after the reform.[39] As highlighted in the histograms we already discussed, there is a substantial rise in stand-alone firms after the reform, and this is mostly driven by new firms. By nine months after the reform, the likelihood of a stand-alone manufacturer has more than tripled relative to pre-reform levels. And while firm characteristics (including vertical integration decisions) can certainly vary based on entry date, the fact that this fraction is quite flat pre-reform and then increases in the period after the reform supports the conclusion that this increase is driven by the reform.

We see a similar pattern for cultivators in the bottom panel of Figure A.4. The fraction of stand-alone cultivators (i.e. those without an operating manufacturer) is quite stable in the pre-reform period and then increases substantially in the period after the reform and the increase is predominantly coming from new firms.[40] By nine months after the reform, the

[39] We define a cultivator as in operation after the first planting date and before the last planting date because this measure will respond to the tax reform quickly.

[40] We define a cultivator as 'stand-alone' analogously to our definition of stand-alone manufacturers. A cultivator is stand-alone either if they have never been associated with a manufacturer that makes a sale to a retailer or for all weeks after the last sale to a manufacturer made by a retailer. Because manufacturers do not always sell every week, it is possible to see a response to the reform in the period immediately preceding the reform if what became the last sale by the manufacturer occurred several weeks prior to the reform.

likelihood of a stand-alone cultivator has more than doubled relative to pre-reform levels.

To summarize, the larger long-run effect we estimate is driven at least in part by some firms switching from one corner solution to another. As our model forecloses on corner solutions, it is perhaps unsurprising that the point estimates of Table 5 are larger than the change predicted by our model.

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

Table 6 reports estimates of Equation (5) with the log of plantings as the outcome variable. Our baseline specification is reported in Column (1). The relevant data is plotted in the top panel of Figure 5. 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. We interpret the level of production prior to the reform as being inefficiently low based on output per unit of capital – firms had the physical capacity to increase their production pre-reform, but did not do it. As soon as the reform occurred, output increased before there was any time to increase their capital. The size of the response could be influenced by the general market expansion identified in Figures A.3 and A.4; however this

effect is likely small because 88 percent of the non-vertical weight sold to retailers from the cultivators in this analysis in the six months after the reform went to manufacturers that were in operation pre-reform.

One may worry that our planting estimates are driven by cyclical or secular trends. These concerns are partially addressed by the placebo estimates one year later provided underneath the main estimates in Table 6, which are approximately zero. Additionally, Figure 5 indicates a clear, permanent shift in plantings immediately following the reform. We also repeat the permutation test that we used for our vertical integration analysis in Figure 6. The gray dashed line in the permutation test figure marks our estimated effect; this line is more than double the next most extreme estimate. We have 50 observations in our permutation test, so the implied p-value is 0.02 (=1/50).[41]

In Columns (2) - (4) we separately examine firms that stand to gain the most from the removal of the GRT structure relative to other firms. Based on the calibrated model that matches our empirical setting, we expect the most salient responses will come from firms that participate in the sell-side of the cultivator-manufacturer market relative to those on the buy-side or firms whose behavior is unaffected by the change in structure (those that don't participate in the cultivator-manufacturer market at all). This is born out in the data. The bottom panel of Figure 5 illustrates the response of cultivating firms on the sell-side relative to all other firms. Before the tax reform, cultivators who participate in

[41]For this permutation test, we begin our placebo estimation in mid-March 2015. Before this, the smaller sample size, driven by many firms still opening and ramping up their production, makes the estimates substantially more noisy and thus a less good comparison to the period of the reform. We use all weeks except the eight weeks surrounding the tax reform and eight weeks surrounding an annual planting slowdown in November and December. However, if we did include earlier periods and weeks surrounding the slowdown, most estimates are still near zero, and the resulting permutation test p-value would be smaller. Hence we are reporting a more informative figure, but also a more conservative permutation test p-value.

the sell-side of the cultivator-manufacturer 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 firms that participate on the sell-side after the elimination of the GRT have less (or no) manufacturer capacity relative to cultivator capacity. The cost disadvantage created by the GRT prevented these firms from selling their plant material to other firms and therefore these firms' production was at an inefficiently low level. This graphical evidence is confirmed by the estimates reported in Column (2) of Table 6. Firms on the sell-side of the cultivator-manufacturer 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 smaller, statistically insignificant amount.[42]

Table 6 Column (3) examines the buy-side of the cultivator-manufacturer market – how do cultivator-manufacturer pairs that purchase plants in the cultivator-manufacturer market after the reform adjust their plantings? 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 effect is small. 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 6 Column (4), it is not

[42]In the two months prior to the reform, firms with non-vertical production after the reform had an average of about 600 plants in the ground that would eventually be harvested, while those without had about 950 plants in the ground. The cumulative effect of the persistent increase in plantings after the reform that we estimate here was to increase the plants in the ground at any given time for those with non-vertical production post-reform to approximately the same level as those without (just shy of 1000 plants in the ground) by six months after the reform.

surprising that the estimates on the interaction terms are similar to when they were examined separately.

The point estimate for the sell-side in Table 6 Column (4) is quite close (within 1%) to the prediction of our model. The point estimate for the buy-side differs slightly, though the confidence interval easily overlaps.

An additional advantage of isolating the responses of sell- and buy-side participants in the cultivator-manufacturer market is that this analysis becomes a difference-in-RDIT analysis with those that don't participate in the cultivator-manufacturer market after the reform acting as a comparison group. This allows us to confirm that our estimates are not affected by other contemporaneous events of factors that influenced planting behavior that would affect all firms in similar ways. Likely, the largest specific concern 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[43]; the fact that Table 6 Columns (2) -(4) find that production significantly expands only for firms that participate on the sell-side of the market (as predicted by theory given our empirical setting) provides evidence that this concern is not the primary mechanism behind the results in this section. To the extent that we are, nevertheless, concerned about residual contamination in the overall estimates and we want an estimate for the entire market, we can apply the estimate in Table 6 Column (2) to the 48 percent of the market that participates on the sell-side of the cultivator-manufacturer market post-reform, and impose no response for the

[43]We have confirmed that the increase in after-tax prices that we document in (Hansen et al., 2020) is approximately the same for both those that do and do not participate in the sell-side of the market post-reform.

other 52 percent of the market. This implies a market-level estimate of 14.5 percent. These calculations further highlight that, as we think about generalizing this estimate to other settings, two parameters are important: (1) the increase in productivity for all firms that are constrained by the tax, and (2) the fraction of the market constrained by the tax. Both of these plausibly vary across settings.

In Table 6, Columns (5) and (6) we consider an alternative dependent variable: the log count of the number of strains planted. The number of strains increases by 11 percent after the reform and this estimate is statistically significant at the 1 percent level. Column (6) demonstrates that this effect is concentrated among firms that participate in the sell-side of the cultivator-manufacturer market post-reform. These estimates imply that these firms are not just expanding their production of existing varieties, but actually increasing the number of varieties that they plant each week in response to the tax reform.

Table 7 and the bottom panel of Figure 6 provide additional robustness checks for our baseline estimate, which are analogous to those we conducted for our analysis of the manufacturer market in Section 6.1. Column (1) repeats the estimate from Column (1) in Table 6. 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.[44] Column (8) adds two lags of the dependent variable. The estimates are

[44] As above, we find similar estimates if we drop firms with any violation and we have confirmed that

highly robust to these different specifications. In the bottom panel of Figure 6, we consider a wide range of different bandwidths and confirm that our estimates are robust to this choice.

7 Welfare

The results of the previous sections align relatively closely with the predictions of our model. A natural question is the extent to which the changes we document in vertical integration and output affect welfare. As we explicitly model demand and cost functions, welfare is well-defined. First, we note that as we do not model retail markups, our model predicts a fall in tax revenue, although we find in related work that the change was ultimately approximately revenue neutral (Hansen et al., 2021). We therefore define a total surplus concept as the sum of consumer surplus and producer variable profits.

Intuitively, a firm facing a high marginal cost for a particular activity will seek to 'outsource' that activity to a firm with a lower marginal cost for that same activity. This implies that firms will specialize according to their comparative advantage, with the degree of specialization depending on the cost function; since we calibrate quadratic cost functions, the specialization is incomplete (both firms engage in both activities). The GRT drives a wedge into this specialization: a firm will conduct more of its higher relative cost activity (and less of its lower relative cost activity) than is optimal. The extent of the deadweight loss (and therefore the extent of the welfare gains that can be generated by the removal of the GRT) largely depends both on the relative costs within each firm (i.e. the extent to which the firms would seek to specialize with no distortions), the relative costs across firms, and the demand

there are no significant changes in the audit or violation rate for cultivators after the tax reform.

elasticity (i.e. the extent to which the effective increase in costs decreases total production).

At the calibrated parameters of Table 3, the removal of the GRT (and the increase in the retail excise tax rate) leads to an increase in total surplus of 2.69%. Of that, 42.5% is captured by firms in the form of higher profits, while 57.5% is passed-through to consumers. In Figure 7, we explore how these predictions change as we change each of the 6 non-level model parameters by up to 10% of their calibrated values (with the exception of the conduct parameter for the final goods market, for which we explore the range from 0 to 1) while holding the other parameters constant.

Our welfare results are most sensitive to changes in the demand elasticity: an increase in the demand elasticity of 10% (to 1.286) increases our predicted total surplus change to 4.62% and decreases the share that accrues to consumers to 52.3%. Increases in the final goods market conduct parameter (i.e. increases in the market power of manufacturers in their transactions with retailers) increases the gains from removing a GRT while decreasing the share that passes through to consumers. While magnitudes differ, welfare is generally more sensitive to changes in firm 1's costs than to changes in firm 2's costs. Intuitively, firm 1 has higher costs for both activities. Note, however, that as the cost to firm 1 of processing cannabis (the highest calibrated cost) increases, the available welfare gains decrease. This illustrates the tension mentioned above: though an increase in this cost increases the gains from specialization, it decreases the gains from changes in total production. To the extent that this cost represents a capacity constraint on the part of a smaller firm, an increase in the cost implies that the constraint is tighter and thus removing the GRT leads to a smaller increase in surplus. Finally, we note that changes in the cost parameters that generate larger increases in the welfare effects of the GRT also generate smaller levels of vertical

integration pre-reform. We conclude that GRTs may have larger impacts on industries with lower baseline levels of vertical integration relative to this industry.

8 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 estimate the elasticity of vertical integration with respect to the intermediate good net-of-tax rate is -0.15 in the short run, and is about twice as large in the long-run. While we are unaware of any other direct estimates of this elasticity, these results can be put into the context of other vertical integration responses. Egger & Seidel (2013) examine the effect of differences in corporate taxes across countries on vertical integration; i.e. in the presence of high tax differentials, vertically integrated firms can shift profits into a low-tax jurisdiction

through transfer pricing. They estimate that an increase in the tax gap between the US and the average host country of 3.1% increases intra-firm trade flows by 5.5% or, in other words, that the tax gap elasticity of vertical integration is 1.77. Alfaro et al. (2016) hypothesize that higher prices in product markets should induce vertical integration and estimate a range of average price elasticities of vertical integration between 0.4 and 2. We conclude that vertical integration is relatively inelastic in this setting.

Our findings are consistent with a simple model of a supply chain with heterogeneous firms. Indeed, after calibrating our model with pre-reform data, our predicted post-reform outcomes are remarkably close to our estimates. The key mechanism is heterogeneous, non-linear (and increasing) marginal costs across firms and activities within the supply chain, which, in our cannabis context, we interpret as capacity constraints. At our calibrated parameters, the 'smaller' firm (i.e. the firm with lower cultivation and processing quantities) faces higher marginal costs for processing than for cultivation. In other words, the firm has 'excess capacity' for cultivation relative to processing, yet the friction of the GRT implies that even if the other firm has some degree of processing capacity available (indeed, the 'larger' firm faces higher costs for cultivation than for processing), the large firm is unwilling to buy much of the small firm's cultivated material. After the removal of the GRT, the small firm is able to expand it's cultivation operations. This leads to modest increases in welfare, the majority of which are passed-through to consumers (as opposed to being captured by firms). The magnitude of welfare changes is primarily driven by the price elasticity of demand: if demand was more elastic, holding everything else constant, we would estimate greater welfare gains from eliminating the GRT.

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

on all 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; we draw this conclusion from exploring the relationship between vertical integration and welfare changes in our model. 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.

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.

Desai, M., Foley, C. F., & Hines, J. R. (2004). The costs of shared ownership: Evidence from international joint ventures. *Journal of Financial Economics*, *73*, 323–374.

Diamond, P. A. & Mirrlees, J. A. (1971). Optimal taxation and public production I: Production efficiency. *The American Economic Review*, *61*(1), 8–27.

Egger, P. & Seidel, T. (2013). Corporate taxes and intra-firm trade. *European Economic Review*, *63*, 225–242.

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.

Grubert, H. (2003). Intangible income, intercompany transactions, income shifting, and the choice of location. *National Tax Journal*, *56*(1), 221–242.

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*, *73*(2), 511–544.

Hansen, B., Miller, K., & Weber, C. (2021). Getting into the weeds of tax invariance. Technical report, National Bureau of Economic Research.

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*.

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*, *43*.

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.

Weyl, E. G. & Fabinger, M. (2013). Pass-through as an economic tool: Principles of incidence under imperfect competition. *Journal of Political Economy*, *121*(3), 528–583.

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	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.
Tier 1			
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
Tier 2			
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
Tier 3			
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: Calibration summary

(a) Parameters

Variable	Description	Value
ϵ_D	Demand elasticity	1.1689
k	Demand level	102.129
β_{g1}	Cost to firm 1 of cultivating cannabis	0.1106
β_{g2}	Cost to firm 2 of cultivating cannabis	0.0616
β_{p1}	Cost to firm 1 of processing cannabis	0.1188
β_{p2}	Cost to firm 2 of processing cannabis	0.0468
θ_p	Conduct parameter in final goods market	0.0417

ϵ_D and k are calibrated from Hansen et al. (2020). The other parameters are calibrated using the below moments.

(b) Pre-Reform Moments

Moment	Data	Model
Fraction vertically integrated	0.955	0.947
Manufacturer to retailer price	3.60	3.60
Cultivator to manufacturer price	1.53	1.53
Firm 1 / Firm 2 cultivation ratio	0.413	0.418
Firm 1 / Firm 2 processing ratio	0.537	0.534

See text for detailed description of moments.

(c) Post-reform predicted outcomes

Outcome	Prediction
Change in vertical integration	-0.022
Change in sell-side plantings	0.306
Change in buy-side plantings	0.018

We model the reform by removing the gross receipts tax and increasing the final goods tax to 37%.

Table 4: Short-Run 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**	−0.023**
	(0.019)	(0.020)	(0.022)	(0.021)	(0.027)	(0.021)	(0.022)	(0.010)
Vertical_{t-1}								0.094***
								(0.008)
Vertical_{t-2}								0.054***
								(0.010)
Placebo Tax Reform	0.007	0.007	0.014	0.010	−0.001	−0.000	0.005	0.011
	(0.012)	(0.012)	(0.015)	(0.013)	(0.018)	(0.012)	(0.014)	(0.007)
Vertical_{t-1}								0.094***
								(0.007)
Vertical_{t-2}								0.060***
								(0.007)
Observations	4,271	4,271	4,271	3947	3,788	3,503	3,279	4,271
Manufacturer 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 Manufacturer Firms	437	437	437	437	437	397	366	437
Regression	Non-linear	OLS	Non-linear	Non-linear	Non-linear	Non-linear	Non-linear	Non-linear
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
Manufacturers 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. The $TaxReform$ coefficients are estimated from equation (5). The outcome is the fraction of vertical transactions for each firm-week. Other variables in equation (5) 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. The correlated random effects included here are equivalent to manufacturer location fixed effects in a linear model.

Table 5: Long-Run Vertical Integration Response

	(1)	(2)
Tax Reform	−0.041*	−0.041*
	(0.024)	(0.023)
New Firms	−0.042*	−0.043
	(0.023)	(0.030)
Placebo Tax Reform	0.021	0.026
	(0.026)	(0.025)
Placebo New Firms	0.003	−0.007
	(0.022)	(0.030)
Observations	17,324	17,324
Manufacturer Firms	562	562
Placebo Observations	9,901	9,901
Placebo Manufacturer Firms	285	285
Correlated Random Effects	No	No
Covariates	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 $TaxReform \; NewFirm$ from equation (7). Other variables in that equation are included, but not reported. The bandwidth is 16 weeks pre-reform and 52 weeks post-reform. The estimates are mean marginal effects from a fractional probit regression for the outcome fraction of non-vertically integrated transactions. The top two rows are the estimates from the tax reform. The bottom two rows 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 6: Production Responses

	Plantings				Strains Planted	
	(1)	(2)	(3)	(4)	(5)	(6)
Tax Reform	0.229***	0.065	0.219**	0.040	0.110***	0.037
	(0.075)	(0.084)	(0.088)	(0.093)	(0.031)	(0.042)
Sell-Side x Tax Reform		0.303***		0.312***		0.129***
		(0.112)		(0.111)		(0.045)
Buy-Side x Tax Reform			0.014	0.041		0.031
			(0.116)	(0.113)		(0.046)
Placebo Reform	0.018	0.004	0.035	0.026	0.002	0.021
	(0.055)	(0.071)	(0.059)	(0.072)	(0.021)	(0.030)
Sell-Side x Placebo Reform		0.019		0.008		−0.009
		(0.084)		(0.082)		(0.031)
Buy-Side x Placebo Reform			−0.056	−0.056		−0.055
			(0.093)	(0.091)		(0.036)
Tax Reform Observations	10,526	10,526	10,526	10,526	10,526	10,526
Tax Reform Cultivator Firms	178	178	178	178	178	178
Placebo Observations	17,115	17,115	17,115	17,115	17,115	17,115
Placebo Cultivator Firms	282	282	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 Columns (1) and (5) are the coefficients on $TaxReform$ from equation (5). The outcome is the log of plantings in the first four columns and the log of the number of strains in the last two columns. Other variables in that equation are included, but not reported. Columns (2) - (4) and (6) interact the tax reform variable with measures for whether the firm bought ('buy-side') or sold ('sell-side') 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 7: Production Planting Response Robustness Checks

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Tax Reform	0.229***	0.224***	0.230***	0.202**	0.303***	0.290***	0.192**	0.153***
	(0.075)	(0.074)	(0.073)	(0.083)	(0.090)	(0.094)	(0.077)	(0.050)
Log(Plantings$_{t-1}$)								0.207***
								(0.015)
Log(Plantings$_{t-2}$)								0.212***
								(0.015)
Observations	10,526	10,526	10,526	10,170	9,458	6,300	9,013	10,526
Cultivator Firms	178	178	178	178	178	100	153	178
Regression	Non-linear	Log OLS	Non-linear	Non-linear	Non-linear	Non-linear	Non-linear	Non-linear
Correlated Random Effects?	No	No	Yes	No	No	No	No	No
Pre-Reform Weeks Excluded	0	0	0	2	0	0	0	0
Post-Reform Weeks Excluded	0	0	0	0	6	0	0	0
Cultivators 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. The $TaxReform$ coefficients are estimated from equation (5). The outcome is the log of plantings. Other variables in equation (5) are included, but not reported. The non-linear estimates are mean marginal effects from a linear hurdle model. The correlated random effects considered in this table are equivalent to cultivator fixed effects in a linear model.

Figures

Figure 1: Vertical Integration Response

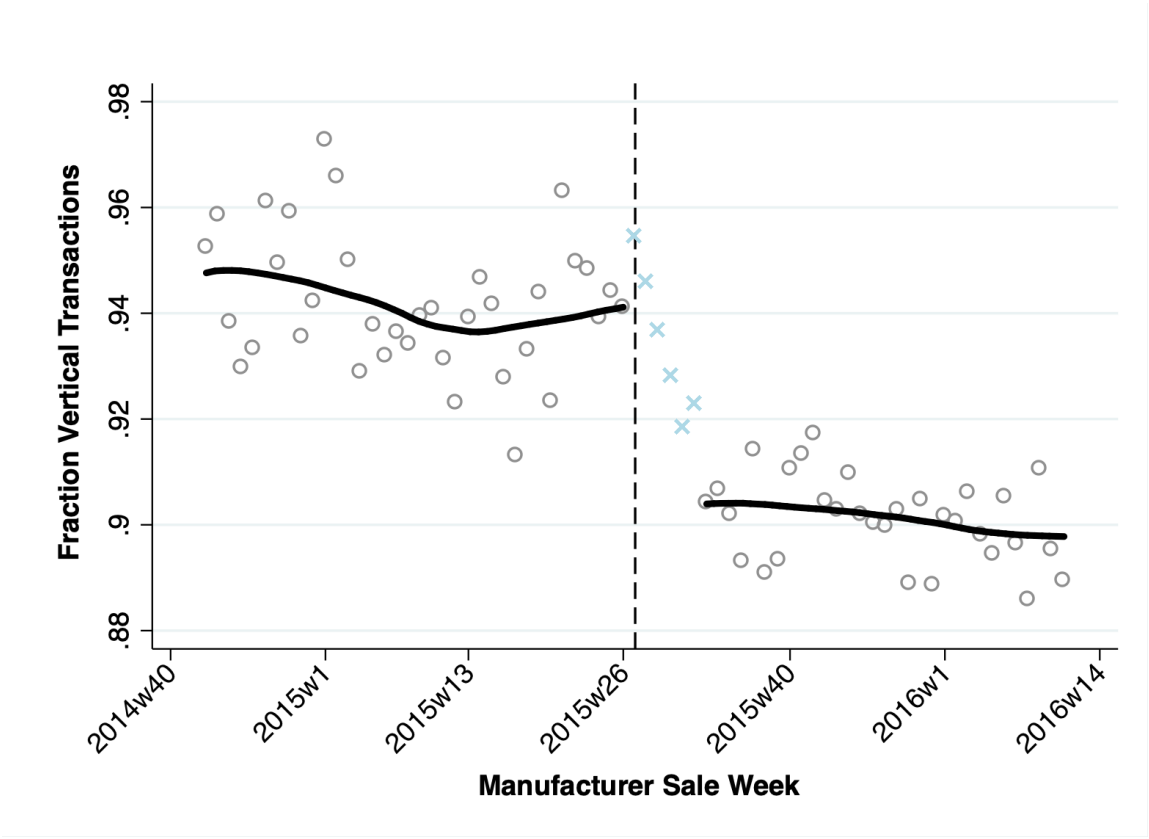

This figure plots the outcome for Table 4 Column (1). The hollow circles represent the raw average of the dependent variable by week. The blue Xs for the six weeks after the tax reform indicate the transition period that we exclude from our main regressions. The solid line is a fourth-order local polynomial 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 Permutation Test & Bandwidth Choice

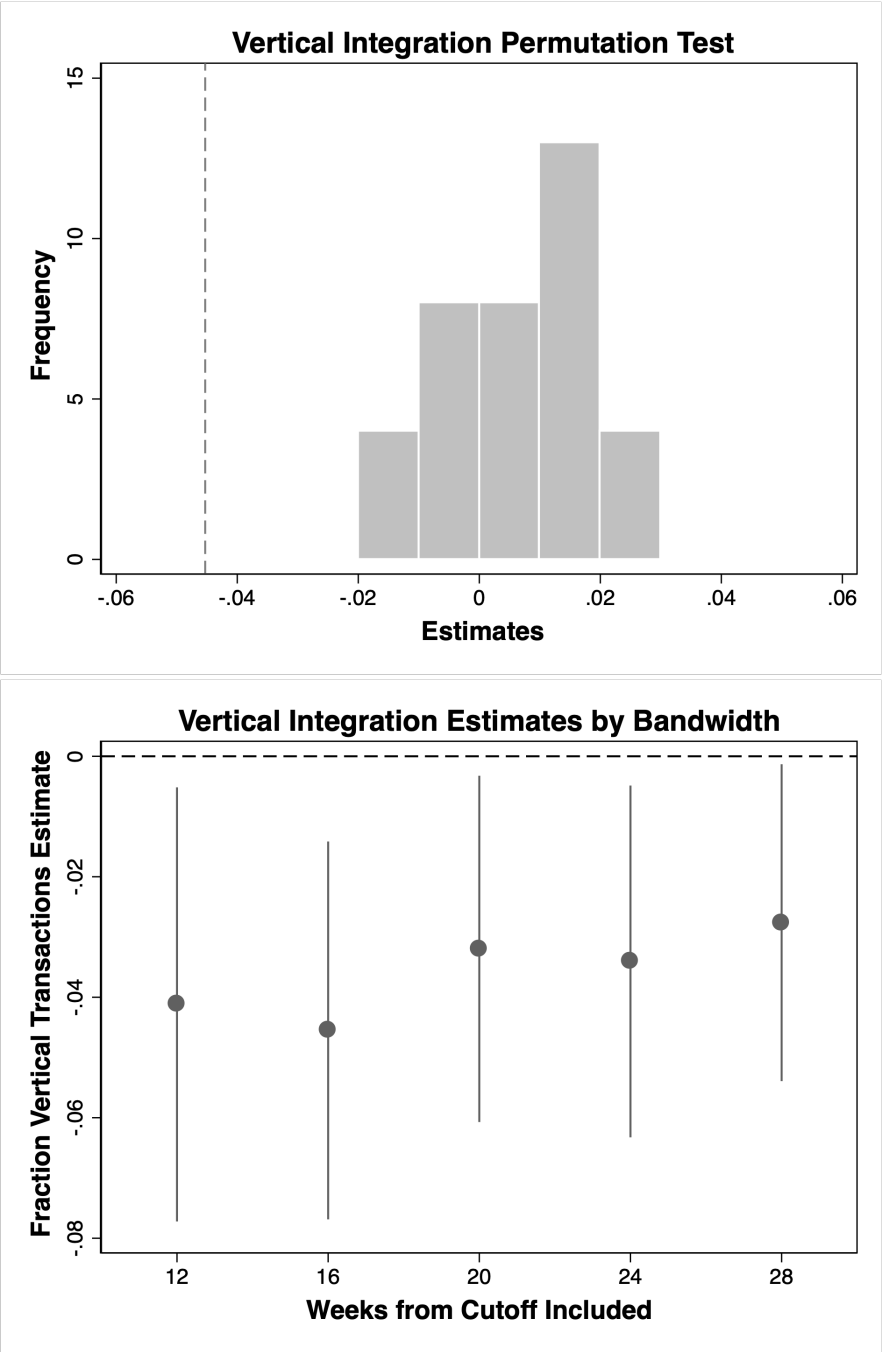

The top panel of this figure plots placebo estimates of of the coefficient on $TaxReform$ in equation (5) that is found in Table 4 Column (1), by reassigning treatment to each week in our data, except those within six weeks of the tax reform, Black Friday, and 4/20. We begin our placebo estimation in mid-February 2015 once there are at least 2,000 observations in our analysis sample. Before this, the smaller sample size, driven by many firms still opening and ramping up their production, makes the estimates substantially more noisy and thus a less good comparison to the period of the reform. However, if we did include earlier periods (and periods we've excluded surrounding holidays), most estimates are still near zero (and our reported estimate remains largest by a wide margin), and the resulting permutation test p-value would be smaller. Hence we are reporting a more informative, but also more conservative, permutation test. The gray dashed line marks our estimated effect; this line is more than double the next most extreme estimate. We have 45 observations in our permutation test, so the implied p-value is 0.022 (=1/45). The bottom panel of this figure plots varied bandwidths for the estimate of the coefficient on $TaxReform$ in equation (5) that is found in Table 4 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

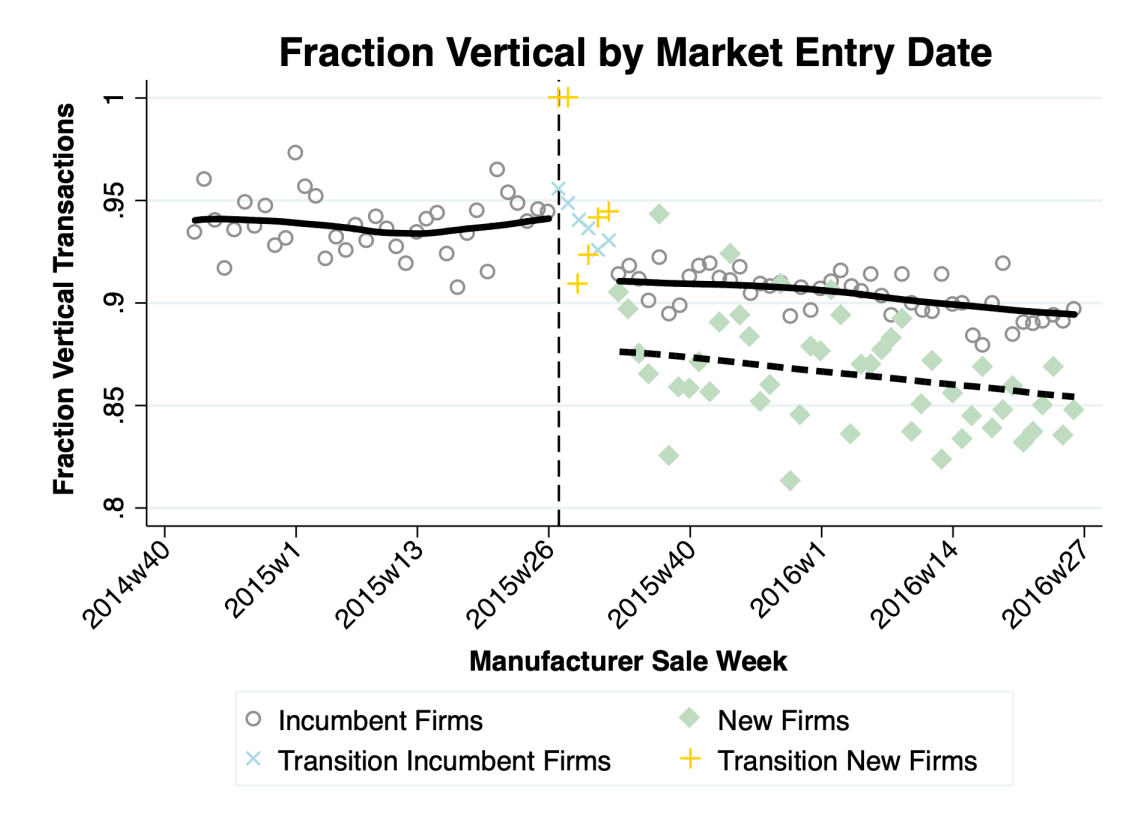

This figure plots the outcomes for Table 5 Columns (1). The hollow circles represent the raw average dependent variable by week for firms that entered before the tax reform (except the six transition weeks post-reform that are excluded from our regression analysis, which are marked by blue X's). The 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 week transition window). The dashed vertical line marks the week of the tax reform.

Figure 4: Vertical Integration Response Histograms

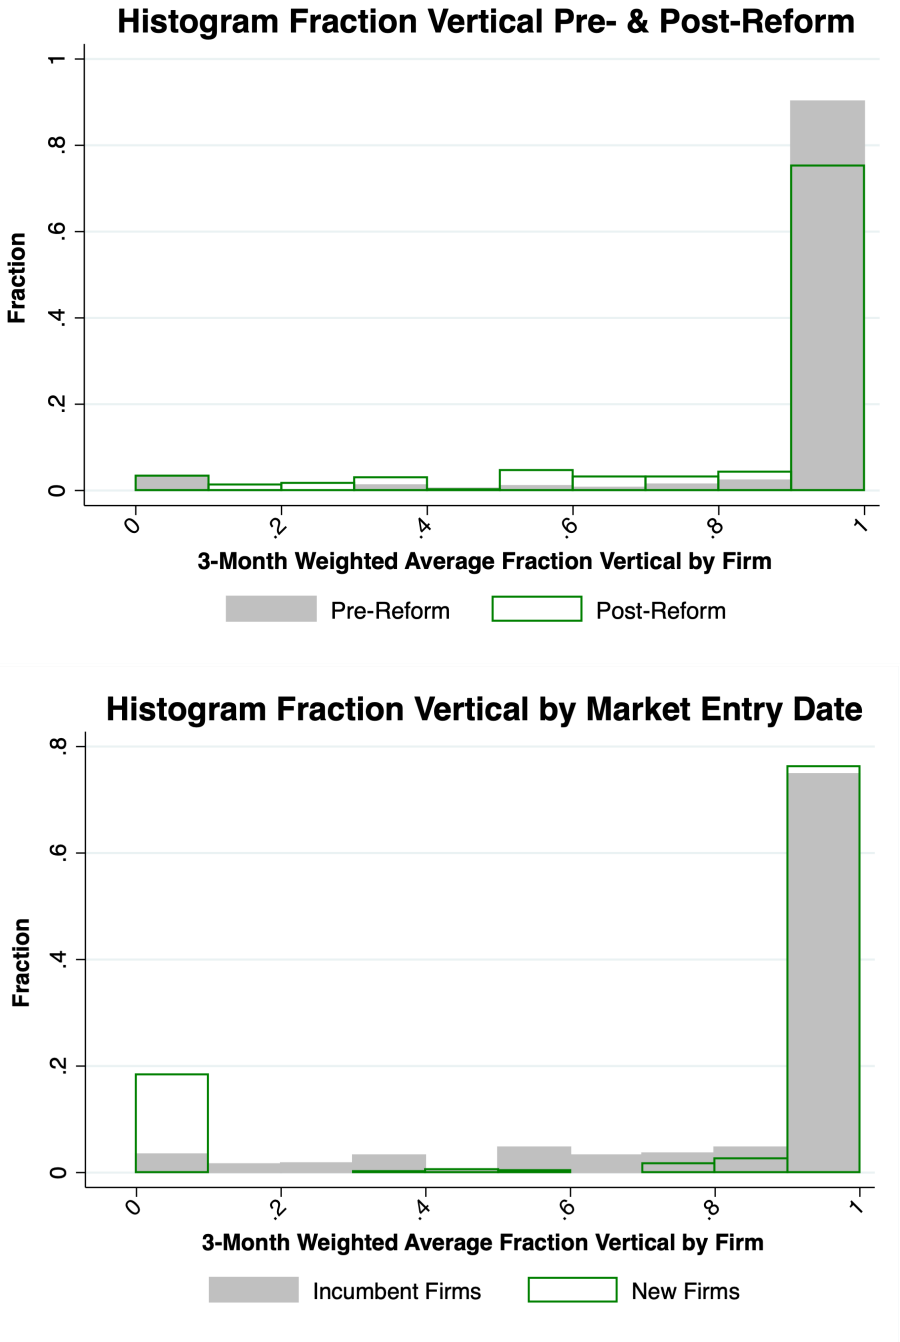

The top panel is a histogram of the fraction of vertically integrated sales by firm for three months before (shaded gray) and after the reform (hollow green), for all firms who were already open prior to the reform (incumbent firms). We take a three-month average of the fraction of vertically integrated activity (weighted by sales). The pre-reform time-period is the 13 weeks (approximately 3 months) immediately prior to the reform. The post-reform period begins six months after the reform, and is again 13 weeks in length. The results are similar for a post-reform period immediately after the reform, but this period further after the reform is chosen to highlight longer-run decisions. The bottom panel replicates the histogram from the top panel for pre-existing firms in the post-reform period and adds a histogram for the same time period for all firms who enter post-reform (new firms) who have entered the manufacturer market by the first week of the three-month average.

Figure 5: Production Response

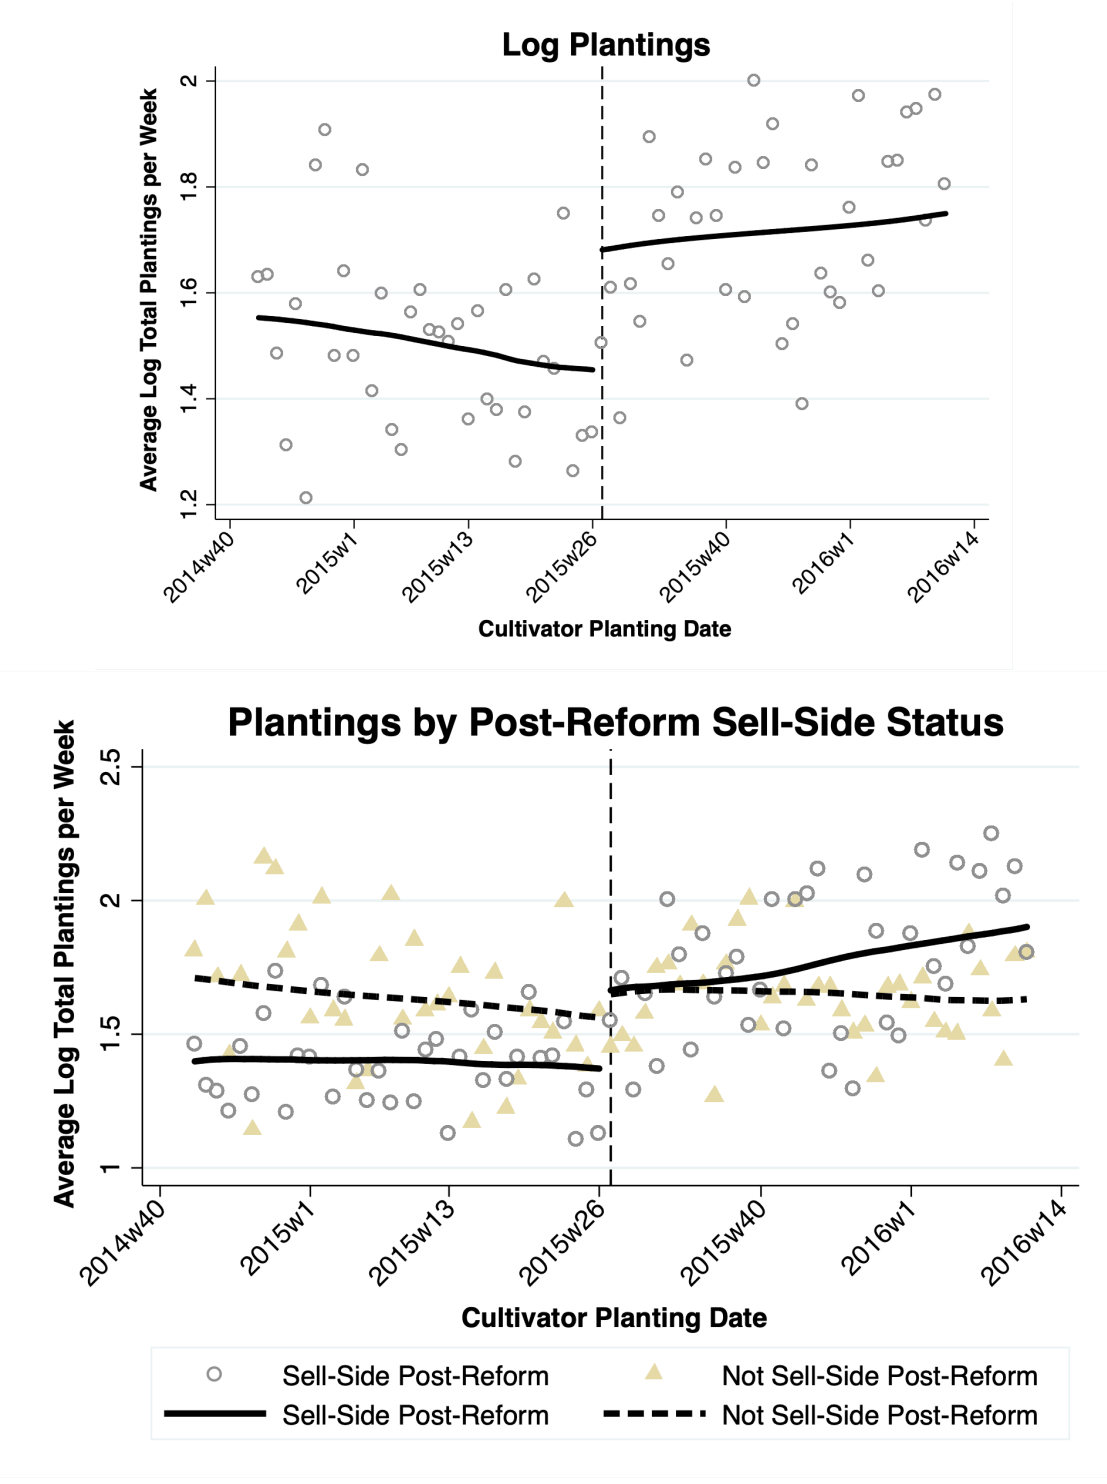

These figures plot the outcome for Table 6 Columns (1) and (2). In the top panel, the hollow circles represent the raw average weekly log plantings. The solid line is a fourth-order local polynomial plot of the raw data. The dashed vertical line marks the week of the tax reform. In the bottom panel, the hollow circles and black solid line represent the weekly log plantings for firms that participate on the sell-side of the 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 on the sell-side of the cultivator-manufacturer market in the year after the tax reform.

Figure 6: Production Permutation Test & Bandwidth Choice

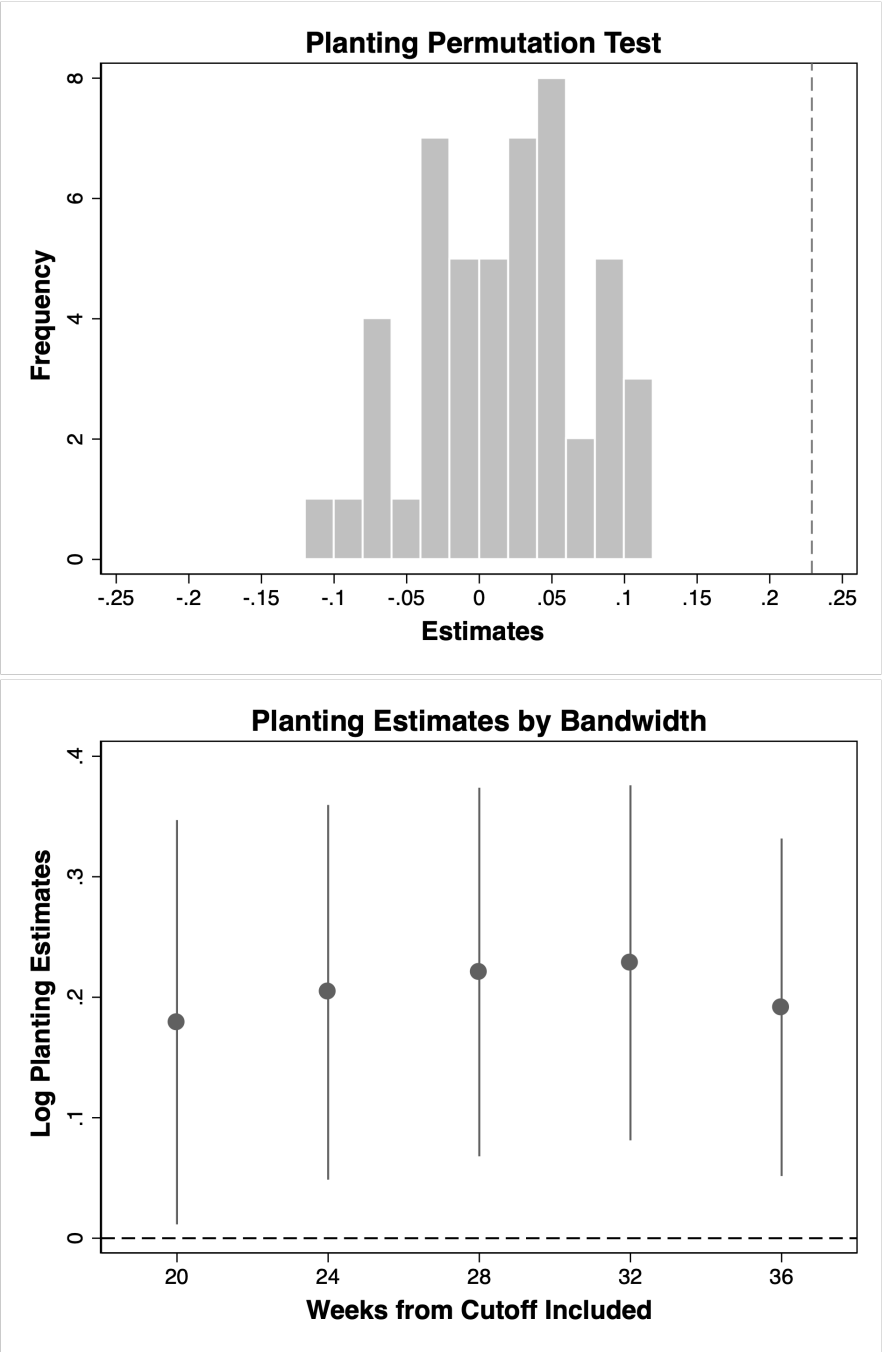

The top panel of this figure plots placebo estimates of of the coefficient on $TaxReform$ in equation (5) that is found in Table 6 Column (1), by reassigning treatment to each week in our data, except those within eight weeks of the tax reform and an annual planting slowdown in November/December. We begin our placebo estimation in mid-March 2015 once there are at least 7,000 observations in our analysis sample. Before this, the smaller sample size, driven by many firms still opening and ramping up their production, makes the estimates substantially more noisy and thus a less good comparison to the period of the reform. However, if we did include earlier periods (and periods we've excluded surrounding the annual slowdown), most estimates are still near zero (and our reported estimate remains largest by a wide margin), and the resulting permutation test p-value would be smaller. Hence we are reporting a more informative, but also more conservative, permutation test. The gray dashed line marks our estimated effect; this line is more than double the next most extreme estimate. We have 50 observations in our permutation test, so the implied p-value is 0.02 (=1/50). The bottom panel of this figure plots varied bandwidths for the estimate of the coefficient on $TaxReform$ in equation (5) that is found in Table 6 Column (1). Recall the main bandwidth is 32 weeks. The dots indicate the point estimates and the lines indicate 90% confidence intervals.

Figure 7: The effect of alternative model parameters on welfare predictions

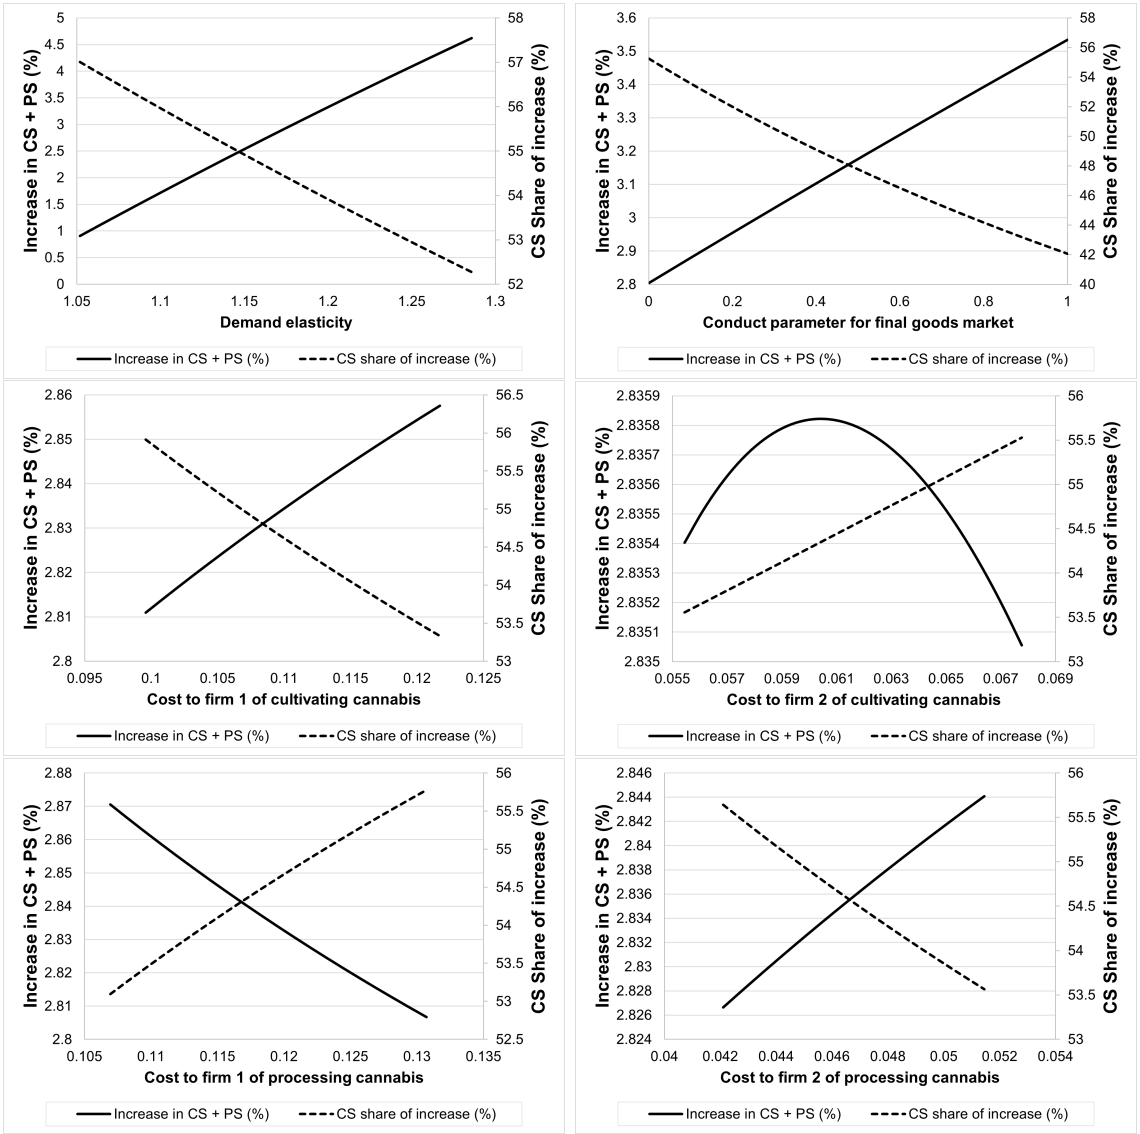

This figure plots the effects of moving from a GRT to a pure excise tax on welfare for different values of the parameters for our model. For each graph, we vary one of the six non-level parameters by up to 10% of its calibrated value as reported in Table 3 (a) and hold the other parameters constant. For the conduct parameter we use the entire range from perfect competition to monopoly, 0 to 1.

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

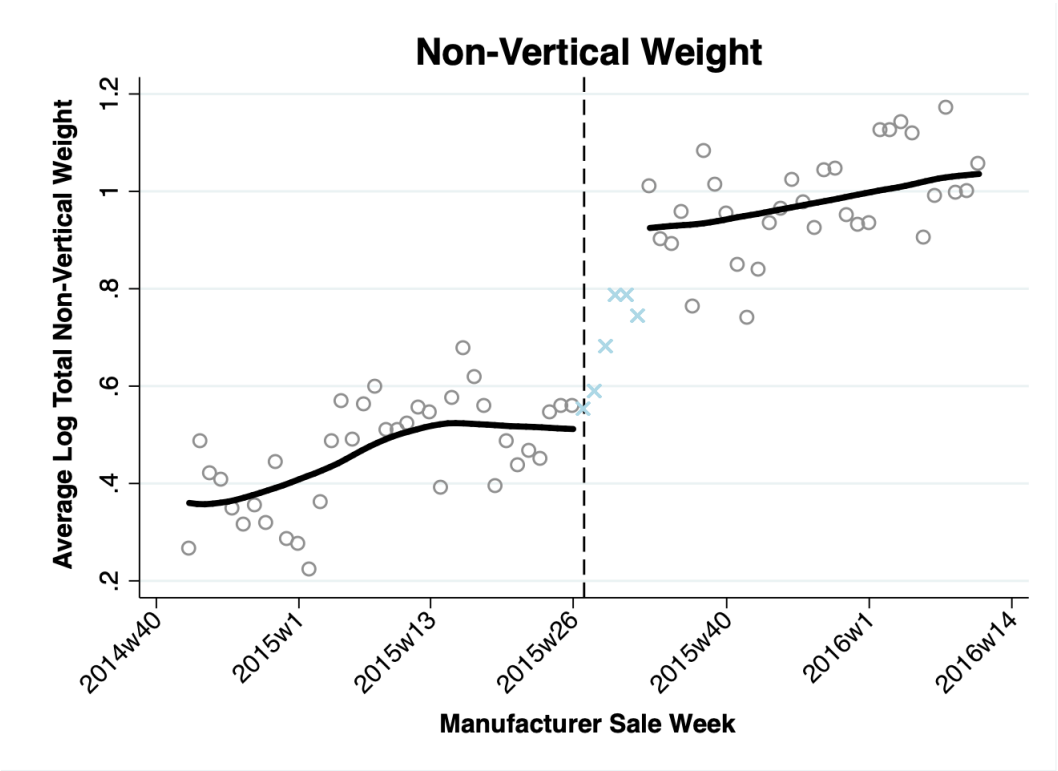

This figure provides an alternative dependent variable – non-vertical weight – to examine vertical integration behavior in response to the elimination of the GRT. The hollow circles represent the raw average weekly dependent variable for each week. The blue X's for the six weeks after the tax change 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 (5), 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

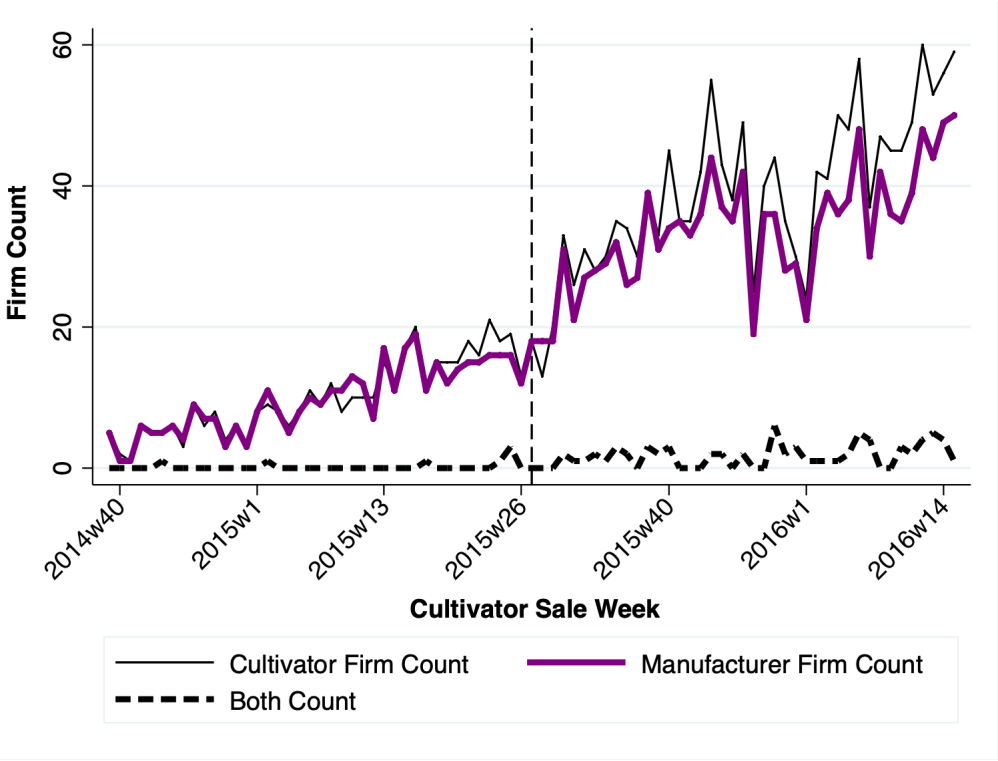

This figure plots the counts of cultivators and manufacturers that participate in the non-vertical cultivator-manufacturer market each week for about 9 months (40 weeks) before and after the reform. The thin black line plots the number of cultivators each week. The thick purple line does the same for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

Figure A.3: Number of Manufacturers

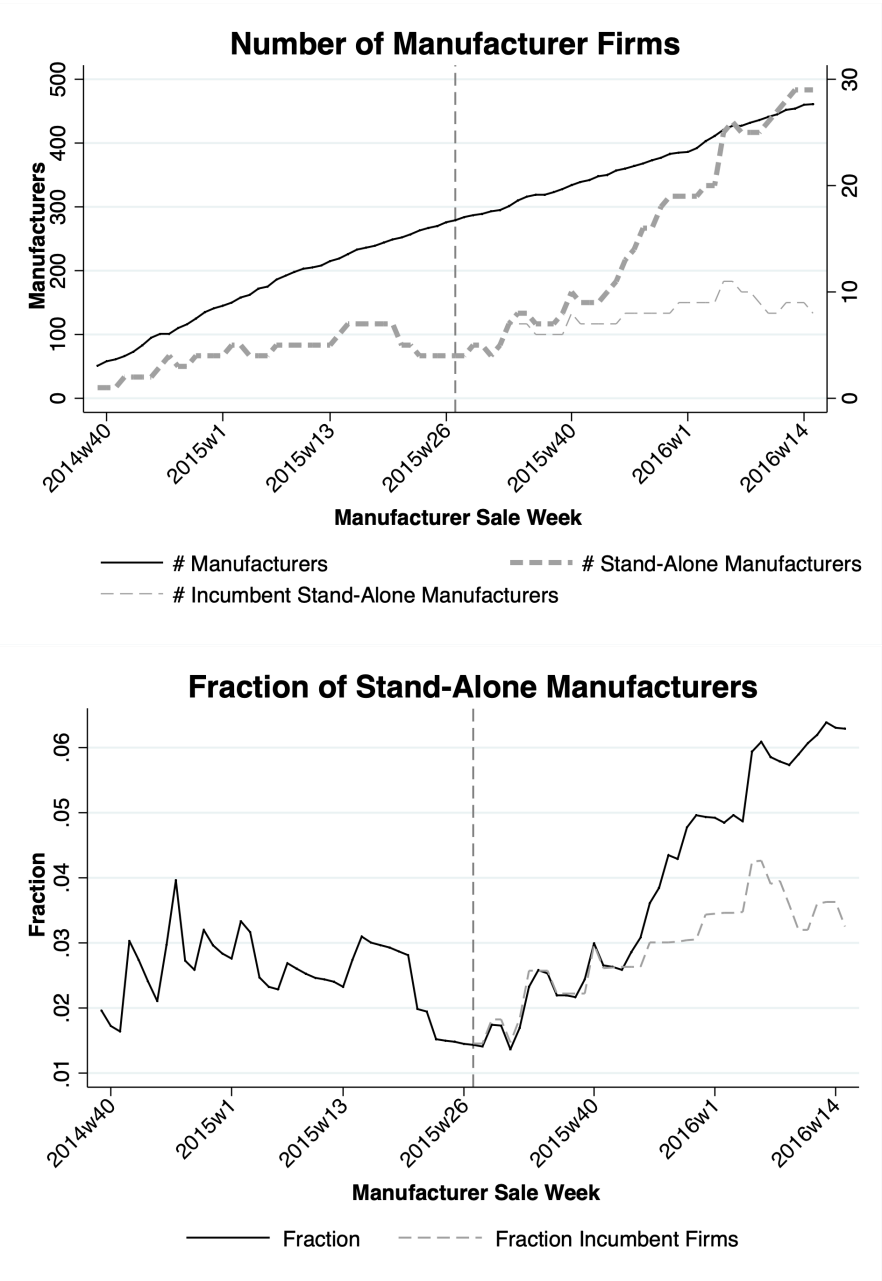

The black solid line in the top figure plots the number of manufacturers in the market each week (as measured by all firms that have made their first sale and have yet to make their last sale) for about 9 months (40 weeks) before and after the reform. The gray dashed line plots the number of stand-alone manufacturers (i.e. those that do not also operate a cultivator, as measured by all firms for whom there is no associated cultivator that has ever planted any cannabis plants to date or for whom the manufacturer did have a cultivator, but the last week that cultivator planted any cannabis has since past. Note: this means we consider the cultivator as shutting down once they stop planting, even though they technically may remain active for a while as they harvest existing plants. This definition allows our measure to respond more immediately to the tax reform and captures all firms who ultimately will no longer have an associated cultivator once harvesting of existing plants is completed). The thin light-gray dashed line marks the number of stand-alone manufacturers across time among firms that had already opened prior to the reform (incumbent firms). The bottom panel plots the fraction of stand-alone manufacturers over time by taking the number of stand-alone manufacturers in the top panel and dividing them by the total number of manufacturers in the top panel. The dashed vertical line in both panels marks the week of the tax reform.

Figure A.4: Number of Cultivators

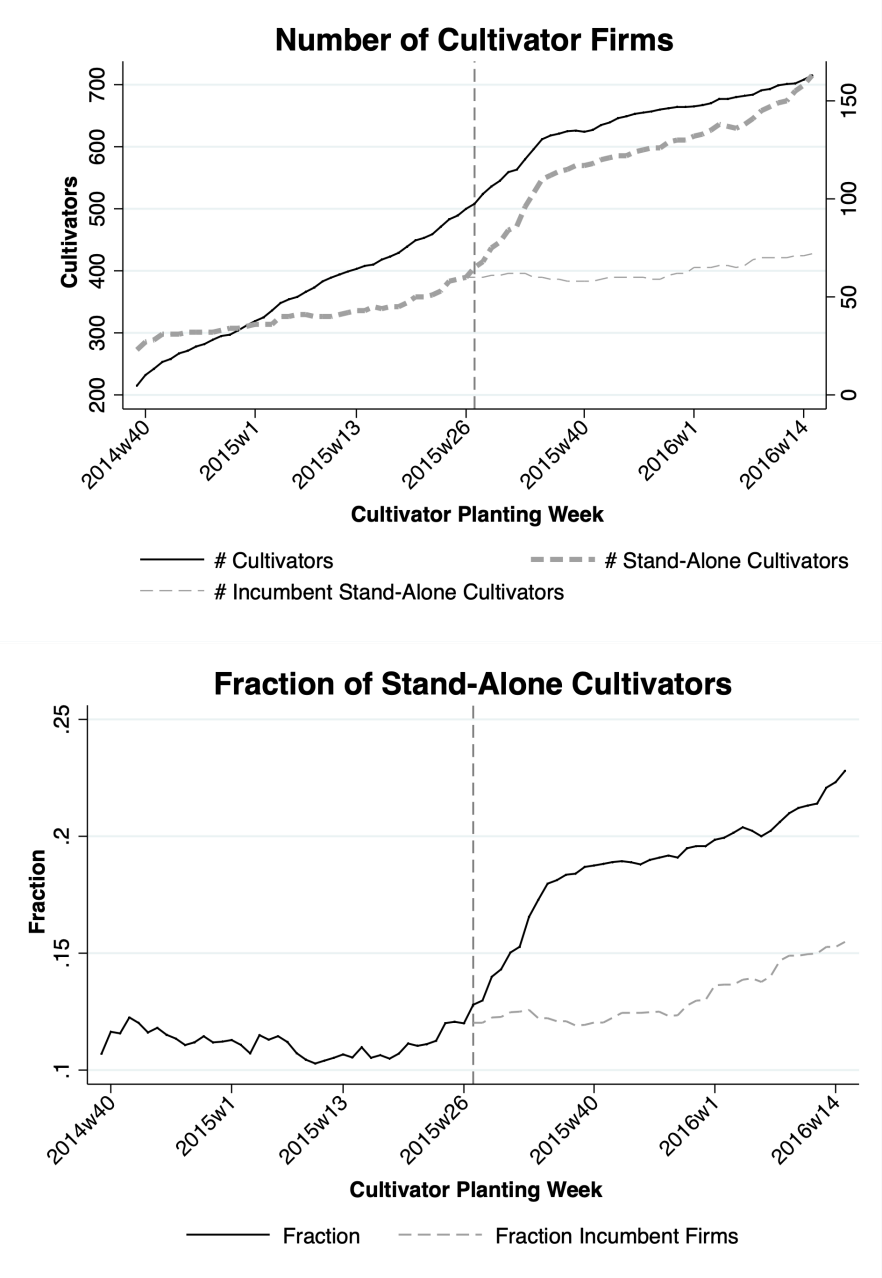

The black solid line in the top figure plots the number of cultivators in the market each week (as measured by all firms that have planted their first plants and have yet to plant their last plants) for about 9 months (40 weeks) before and after the reform. The gray dashed line plots the number of stand-alone cultivators (i.e. those that do not also operate a manufacturer (as measured by all firms for whom there is no associated manufacturer that has ever sold any cannabis to retailers or for whom the cultivator did have an associated manufacturer, but the last week that manufacturer sold any cannabis has since past.). The thin light-gray dashed line marks the number of stand-alone cultivators across time among firms that had already opened prior to the reform (incumbent firms). The bottom panel plots the fraction of stand-alone cultivators (and fraction of incumbent stand-alone cultivators) over time by taking the number of stand-alone cultivators in the top panel and dividing them by the total number of cultivators in the top panel. The dashed vertical line in both panels marks the week of the tax reform.

Figure A.5: Number of Plantings

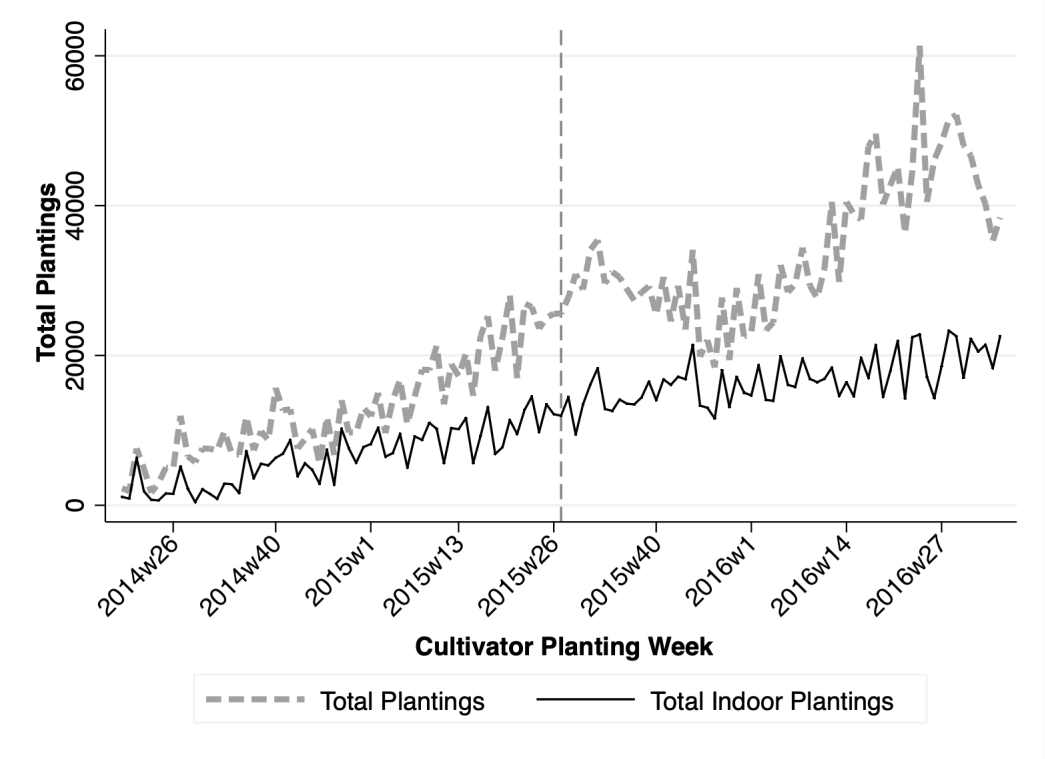

This figure plots the total number of plantings for all firms and for indoor-only firms. Note that in this figure there is no restriction on when firms entered the market, so the composition (and number) of firms shifts over time. So, this figure is more useful for understanding the overall market than it is for understanding firm's responses to the tax reform.