Assessing the Effects of Establishment Entry on Competition and Variable Profits

Joshua Scott[1*] Murad Zeynalli[2]

Eddy Junarsin[1]

[1]Department of Economics, State University of New York at Cortland, PO Box 2000, Cortland, NY 13045

[2]Department of Economics, Miami University, Oxford, OH 45056

Abstract

Does new establishment entry decrease variable profits for existing competitors in a regulated market? This paper studies how retail entry affects incumbent variable profits in Washington's recreational cannabis market. We combine seed-to-sale transaction data with plausibly exogenous variation generated by the state's retail license allocation rules, lottery-based assignment in oversubscribed jurisdictions, and the 2016 expansion of retail license caps. In the baseline specification, we utilize an instrument set built from the log square root of county area and policy-interacted within-county rank measures tied to the 2016 cap expansion. We find that one additional retail establishment in a market reduces an incumbent's monthly variable profits by 4.2% and lowers monthly revenue by 3.6%. The decline operates primarily through lower quantities sold rather than lower prices. The effects are substantially larger in jurisdictions that were already close to their relevant license caps before the 2016 expansion, which indicates that entry matters most where regulation had kept local competition especially scarce. The main findings are also robust to boundary-based market-definition checks and alternative competition measures. These results provide new evidence on the profit effects of entry in a regulated differentiated retail market, while also showing that the external relevance of entry estimates depends importantly on institutional setting when prices do not respond in the way documented in other retail contexts.

Keywords: cannabis, entry, competition, variable profits, market structure

JEL Codes: L11, L13, L81, C26

*Corresponding author: Joshua Scott, joshua.scott03@cortland.edu.

1 Introduction

Since the foundational work on entry and competition by Bresnahan and Reiss (Bresnahan and Reiss, 1991), there have been numerous works assessing the effects of entry on various components or outcomes related to market structure.[1] One particular aspect that has been established is the concept that entry of new firms reduces profits of existing firms (Stevenson and Wolfers, 2020). Despite the strong theoretical foundation, it is generally difficult to uncover causal impacts of entry due to the endogenous nature of entry and the limitation that the data necessary to assess such effects is often unavailable or confidential.

This paper studies how retail entry affects incumbent profits in a regulated retail market. That question is central to industrial organization, but it is difficult to answer credibly because entry is endogenous and direct measures of variable profits are rarely observed. Washington's recreational cannabis market is unusually well suited to the problem because retail licensing rules, lottery-based assignment in oversubscribed jurisdictions, and a later cap expansion generate policy-driven variation in realized local entry.

Washington provides three advantages for this exercise. First, the legal recreational market begins essentially from scratch, so we observe the evolution of entry, sales, and product mix from the beginning of the market. Second, the state's seed-to-sale tracking system links retail transactions to upstream wholesale purchases, which allows us to construct store-level outcomes for revenue, quantity, price, margins, markups, and variable profits. Third, retail entry was shaped by administrative rules that generated plausibly exogenous variation in the number of stores allowed across local markets.

Using these features of the data and policy environment, we estimate the effect of retail entry on incumbent performance with an instrumental-variables design. In the preferred baseline specification, one additional establishment in a Census Designated Place (CDP) reduces an incumbent's monthly variable profits by about 4.2% and lowers monthly revenue by about 3.6%. The revenue effect is driven by lower quantities sold. We do not find a corresponding decline in average retail prices in the preferred baseline design, which points to institutional features of this market—including binding licensing restrictions, residual market power, and pricing practices that appear to respond more strongly to wholesale costs than to local entry. Those conclusions continue to hold when we exclude boundary-adjacent stores, redefine the endogenous competition count to remove boundary stores, and replace the CDP market definition with 10-, 15-, and 20-minute drive-time competitor counts.

The paper contributes to two strands of literature. The first is the broader literature on entry and competition. Experimental and quasi-experimental studies have shown that entry can affect prices, quality,

[1] Examples include studies of how entry affects variable profits and quantity in local hospital markets (Abraham et al., 2007), product variety and welfare in radio broadcasting (Berry and Waldfogel, 1999), and prices and quality in retail markets (Busso and Galiani, 2019).

and market conduct in retail and intermediary markets (Busso and Galiani, 2019; Bergquist and Dinerstein, 2020; Hollenbeck et al., 2024). Direct evidence on the profit consequences of entry remains much more limited because profit measures are typically unavailable. This setting makes it possible to connect entry not only to prices and revenues, but also to variable profits and the margins through which those profits change.

The second contribution is to the growing literature that uses Washington's administrative cannabis data to study taxation, market power, vertical organization, spatial spillovers, and firm scale. Existing work uses these data to study cross-border demand (Hansen et al., 2020), tax incidence (Hansen et al., 2017), vertical integration and tax distortions (Hansen et al., 2022), market power (Hollenbeck and Uetake, 2021), and the returns to operating multiple stores (Hollenbeck and Giroldo, 2022). Relative to that literature, this paper focuses on the incumbent-side consequences of local retail entry itself.

We keep the main estimating framework transparent and modular, while additionally accounting for heterogeneity by baseline competition, heterogeneity by regulatory slack, explicit discussion of the IV estimand, and targeted falsification exercises. This structure is useful not only for internal validity, but also for clarifying external validity. The policy variation identifies a credible local effect of entry in this setting, but the extent to which those magnitudes transport to other markets depends on the institutional environment in which entry occurs (Roe and Just, 2009).

The remainder of the paper proceeds as follows. Section 2 describes the institutional background, the data, and sample construction. Section 3 presents the identification strategy and empirical specification. Section 4 reports the main results. Section 5 concludes and discusses limitations and natural extensions.

2 Institutional Context and Data

Washington legalized recreational cannabis in 2012 and subsequently created the regulatory institutions needed to oversee the new market. Washington's approach differs from that of many other states in two ways that are central to this paper. First, the state restricted retail entry through explicit license caps. Second, the state maintained unusually detailed administrative records from seed to sale and made those records available upon request. The Washington State Liquor and Cannabis Board (WSLCB) also conducts regular compliance visits throughout the supply chain. This section summarizes the institutional setting, the licensing process, and the construction of the estimation sample.

2.1 Initiative 502

Washington legalized recreational cannabis for adults age 21 and over through Initiative 502, passed in November 2012. I-502 gave the Washington State Liquor and Cannabis Board until December 1, 2013 to craft rules governing implementation and regulation. Two central goals of the rules were to regulate the legal market in a way that protected public safety and to reduce the role of illegal drug markets within the state.

I-502 created three tiers in the legal supply chain and established licenses for each tier: producers, processors, and retailers. Vertical integration across the entire supply chain was restricted, although producers and processors were allowed to integrate (State of Washington House of Representatives, 2012). Retailers had to remain separate from the rest of the supply chain. In September 2013, the WSLCB announced that BioTrack THC would provide traceability software for the state (State of Washington Liquor Control Board, 2013a). That software tracks transactions from seed to sale. Combined with the requirement that the supply chain remain within the state, the system generates the detailed data used in this analysis.

2.2 Retail License Allocation

Washington placed an upper limit on the number of retail establishments allowed to operate in the state. The WSLCB relied on BOTEC Analysis Corporation to help determine the cap and allocate licenses across markets (State of Washington Liquor Control Board, 2013b). Licenses were first allotted by county using a non-linear formula designed to "allocate the designated number of stores across counties in a way that minimizes the population-weighted average distance over all the counties" (Caulkins and Dahlkemper, 2013). BOTEC used 2010 NSDUH survey data to estimate the number of users and then used an optimization tool in Excel that yielded a fractional number of stores per county subject to a statewide target of 330. Those allocations were then rounded to whole numbers. If the fractional allocation fell between zero and one, BOTEC rounded up to one.

From there, licenses were allotted to 128 qualified jurisdictions based on their share of county population in the 2010 Census. The qualified jurisdictions were the most populous incorporated cities within each county. Because those allocation rules were set before the market was fully operating, they generate plausibly exogenous differences in the number of stores that jurisdictions were permitted to host, even across places with otherwise similar retail potential. The optimization targeted 330 statewide licenses, but rounding the fractional county allocations to whole numbers produced an implemented initial total of 334 licenses. In 2016, the cap increased to 556 to incorporate medical dispensaries into the legal system (Washington State Liquor and Cannabis Board, 2015). This later expansion provides an additional source of variation in potential competition. Figure 1 maps retail cannabis locations and Census Designated Places in Washington as of October 2017.

2.3 Retail License Lottery

In June 2013, the WSLCB announced that it would begin accepting marijuana-license applications in September of that year.[2] Applications cost $250, and successful applicants paid a $1,000 license fee. Because the number of applications exceeded the number of available retail licenses in many jurisdictions, the WSLCB used a lottery in 75 jurisdictions (Washington State Liquor and Cannabis Board, 2014). The remaining 47 jurisdictions did not require a lottery.

Applications that passed the pre-qualification process became eligible for the lottery. A double-blind lottery was held on April 21–25, 2014, and the WSLCB posted the results shortly thereafter.[3] Each applicant was randomly assigned a number starting at 1. If an applicant's rank was no greater than the number of licenses allotted to that jurisdiction, the applicant became eligible to receive a license. Results were posted in May and licenses were issued no later than July. To limit local dominance by a small number of firms, no firm could hold more than three licenses, and a multi-location licensee could not hold more than 33% of the allotted licenses in any county or city.

2.4 Retail Sales and Seed-to-Sale Monitoring

Retail sales of recreational cannabis began in July 2014. Washington used the BioTrack THC traceability system to follow each item from seed to sale and to enforce the rules established under I-502 and subsequent regulations. The state also sought to remain within the broad federal enforcement guidance described in the August 2013 Cole Memorandum (Cole, 2013). The system was enforced by the state with a particular focus on preventing tax evasion. One important rule is that retailers cannot sell cannabis without manufacturing records. Because each package has a traceable inventory ID, that rule gives retailers a strong incentive to report transactions accurately. Retailers are audited by the WSLCB and face one or more visits per year. Non-compliance is subject to substantial fines (Hansen et al., 2017).

Table 1 reports entry statistics using establishment-level reporting data from the Washington traceability system. The first column indicates the year. Column 2 reports the number of establishments open in the first month of the year, and column 3 reports the number operating in the final month of the year. The first observations are from July 2014, when 17 establishments were operational. Even though retail licenses were capped and 75 jurisdictions had more applicants than available licenses, entry increased only gradually. By June 2015, roughly one year after recreational sales began, only 161 firms were reporting sales. The number of reporting locations did not exceed the initial cap of 334 until December 2016.

[2] License applications officially closed in March 2016.

[3] See the WSLCB press release on lottery results: https://lcb.wa.gov/pressreleases/lottery-results-cannabis-retail-stores-available-wslcb-website.

2.5 Data

2.5.1 Washington Traceability Data

The central feature of the data is Washington's seed-to-sale traceability system for recreational cannabis transactions. These data allow the researcher to observe both retail and wholesale transactions at the individual-transaction level from July 2014 through October 2017. The retail dispensing data alone contain more than 110 million observations over that period. Other states also used BioTrack THC software, but Washington's policy of making the data available is unusual.

Products are categorized as "usable" cannabis, edibles, inhalables, and topicals. Usable cannabis refers to flower intended to be smoked and accounts for roughly three-quarters of all retail transactions in the sample. These products are differentiated by strain and by THC and CBD potency (Hansen et al., 2017). Flower is sold by weight, whereas the other product categories are sold by the unit.

The registration process begins with each harvest. Harvested plants are recorded by inventory lot, and each lot is assumed to be internally homogeneous. Those lots are later split into smaller lots when products are sold to retailers. Retail lots contain many sealed packages of a pre-specified weight or unit size. When those lots are sold, they are logged in the system and assigned a new unique ID. That ID identifies the manufacturer, the product, the size of the lot, and the retailer that receives it (Hansen et al., 2017). We use that information to link wholesale prices to retail sales.

The traceability data also include retailer, processor, and producer information. For retailers, we use address, license number, sales-location ID, and organizational ID in estimation. We combine the traceability data with BOTEC reports on license allocation and lottery jurisdictions in order to construct the instruments used in the analysis.

2.5.2 National Historical Geographic Information System (NHGIS) and U.S. Census Data

For demographic and geographic controls, we use data from IPUMS NHGIS (Manson et al., 2021). The relevant series-5 surveys cover Census tracts, Census Designated Places, and counties. For this study, we use tract, CDP, and county populations.

The geographic boundaries come from the U.S. Census Bureau. We use CDPs to define markets because the jurisdictions that receive licenses generally follow those boundaries. Census tracts and their populations are used to construct population buffers around each establishment. Those buffers enter the estimating equations as controls for local demand conditions.

2.5.3 Constructing the Sample Used in Estimation

Using the unique retail-lot IDs and data on both wholesale and dispensing transactions, we match upstream purchases to downstream retail sales and then aggregate the data to the retailer-month level. Retailer data are matched to establishments using the sales-location ID. We clean the data to remove clear outliers. The maximum legal sale is one ounce, so we drop all retail sales above 28.5 grams.[4] We also drop prices per gram above $40,[5] and we drop prices below $1 to remove product samples that were logged in the point-of-sale system. Retailers that report sales in only one month are also dropped from the sample. Location IDs and addresses are cleaned to ensure that the geographic matches are correct.

For the main analysis, we calculate monthly total variable profit, variable profit for products sold by weight, and variable profit for products sold by the unit. We construct analogous measures for revenue, quantity sold, price, gross margin, and markup. Prices are tax exclusive. Retailer addresses and sales reports are then used to construct the main explanatory variable: the number of establishments in a market each month, where the market is defined as a CDP.

Retailer exit is comparatively limited over the sample horizon. That feature matters for interpretation. The estimates are best understood as showing how continuing incumbents adjust variable profits, revenue, quantities, and prices when local entry increases, rather than as a setting in which entry rapidly induces widespread contemporaneous store closure.

The geocoded store locations are linked to population data at the tract, CDP, and county levels. We also calculate population within 1-, 2-, 3-, 5-, and 7-mile radii around each location. Figure 2 maps Census tracts in Washington and in the bordering portions of Oregon and Idaho together with retail locations. We use the 1-mile and 2-mile radii as demand controls in the preferred regressions. CDP populations are used to construct the rank-based instruments in the preferred design. The same geocoded locations are also used to build the spatial robustness variables introduced later in the paper. In particular, we measure each retailer's distance to the nearest CDP and county boundary, create indicators for stores located within 500 and 1,000 feet of a CDP boundary, and use the ArcGIS origin-destination cost matrix to count active competing retailers within 10-, 15-, and 20-minute drive-time bands. Those variables are used only in the alternative-market and boundary-based robustness exercises and do not change the preferred baseline sample construction.

[4] Approximately 0.16% of observations.
[5] Approximately 0.05% of observations.

3 Empirical Strategy

The combination of retail entry restrictions, the license-allocation formula, the double-blind lottery, and the later cap expansion makes it possible to estimate the effect of entry on store-level outcomes using an instrumental-variables strategy. This section presents the empirical specification, explains the institutional variation used for identification, and clarifies the estimand.

3.1 Estimating Equation

We estimate the effect of entry on store-level outcomes using an instrumental-variables specification:

$$\ln(y_{rt}) = \beta N_{pt} + \gamma \boldsymbol{X}_r + \alpha_p + \xi_m + \lambda_\tau + \epsilon_{rt}, \tag{1}$$

where r indexes retailers, t indexes months, p indexes Census Designated Places, m indexes calendar months of the year, and τ indexes years. The dependent variable is the natural log of variable profit, revenue, quantity sold, or price, depending on the specification. The key explanatory variable, N_{pt}, is the number of retailers reporting sales in CDP p during month t.

The control vector X_r includes the natural log of population within one mile of the retailer, the natural log of population within two miles of the retailer, and the number of establishments in the retailer's organization. The population controls proxy for a store's local demand environment. We include organizational size because previous work shows that multi-store firms earned higher per-store profits over time, partly because they reduced wholesale costs and margins more effectively (Hollenbeck and Giroldo, 2022). We also include an indicator for the post-legalization Oregon border shock documented by Hansen et al. (2020). The indicator equals one for Washington retailers on the Washington-Oregon border beginning in October 2015, when Oregon's recreational legalization began to affect cross-border demand.

Since entry is endogenous, we instrument N_{pt} in the estimating equation. The endogeneity in observed retail entry arises from firms/establishments choosing whether to open in response to expected profitability. The same unobserved local demand, cost, and competitive conditions that affect variable profits can also affect realized entry (Berry, 1992; Angrist and Krueger, 2001). Although Washington used a lottery in oversubscribed jurisdictions, the lottery only ranked pre-qualified applicants within locally capped markets and did not randomize which jurisdictions received more slots, whether local bans or siting restrictions blocked opening, or whether selected applicants actually entered. Therefore, we instrument realized entry with policy-driven allocation variation rather than treat observed entry as exogenous (Washington State Liquor and Cannabis Board, 2014, 2016).

3.2 License Allocation and Instruments

The WSLCB contracted BOTEC Analysis Corporation to provide technical expertise during implementation of I-502. BOTEC's report documents five different methods for allocating retail licenses.[6] The fifth method, which minimizes a proxy for the average distance from a user to the nearest store, was ultimately chosen. BOTEC implemented the method using a non-linear optimization routine in Excel.

The allocation formula can be summarized as

$$\{n_c^*\}_{c=1}^C = \underset{\{n_c>0\}_{c=1}^C}{\arg\min} \sum_{c=1}^{C} \frac{1}{2}\sqrt{\frac{POP_c^2 Area_c}{n_c}}. \tag{2}$$

where n_c is the continuous allocation chosen for county c, POP_c is either county population or the estimated number of users in county c, and $Area_c$ is county area (Caulkins and Dahlkemper, 2013). The equation summarizes BOTEC's distance-minimization objective across counties. BOTEC notes that the procedure is imperfect because population is not uniformly distributed within counties and stores are not placed literally to minimize distance. Even so, the mechanics of the formula and the subsequent rounding generate useful quasi-exogenous variation in the number of licenses assigned across local jurisdictions.

The WSLCB used the BOTEC model and ultimately chose a statewide target of 330 stores when implementing the allocation. After rounding the resulting county allocations, the actual total became 334 licenses. Those licenses were then allocated among cities within counties according to population shares. Seventy-five jurisdictions used a double-blind lottery to distribute licenses to applicants, while 47 did not. Throughout the paper we use the terms *jurisdiction* and *CDP* interchangeably because the allocated jurisdictions largely line up with CDP boundaries.

We use this allocation process to construct instruments for the number of establishments in a market. On December 15, 2015, the WSLCB announced that the retail cap would increase beginning on January 6, 2016 in order to incorporate medical dispensaries into the legal system (Washington State Liquor and Cannabis Board, 2015). The number of available licenses doubled in the ten counties with the highest medical sales, while the remaining eligible counties saw their caps rise by 75%. Skagit and Cowlitz Counties received 100% increases instead of Yakima and Benton because of local bans. Let

$$D_{jt} = NoBan_j \times Post2016_t \times LicenseDouble_j. \tag{3}$$

Here, $NoBan_j$ equals one when jurisdiction j does not prohibit cannabis retailing, $Post2016_t$ equals one in 2016 and later, and $LicenseDouble_j$ equals one for counties and jurisdictions whose retail caps doubled rather

[6] The five methods were simple pro-rating by population, pro-rating by past-month users from the NSDUH, a linear sum with area, a linear sum using the square root of area, and minimization of a proxy for average distance to the nearest store.

than increasing by 75%. The preferred excluded-instrument set therefore consists of the log square root of county area and three policy-interacted allocation measures: within-county population rank, an indicator for the largest jurisdiction in the county, and an indicator for the second-largest jurisdiction in the county, where each rank-based component is interacted with D_{jt}.

The county-area instrument is the natural log of the square root of county area, which is motivated directly by BOTEC's use of county area in the allocation formula. Population rank matters because the jurisdiction with the highest population in a county is mechanically favored in the within-county allocation of licenses, and the same logic applies, although more weakly, to the second-largest jurisdiction. The advantage becomes much smaller after rank 2 because the number of allocated licenses declines rapidly with population rank. County area enters the preferred specification un-interacted because it is the cross-county component of BOTEC's allocation rule. The rank, rank-1, and rank-2 components enter through D_{jt} because the 2016 cap expansion is what turns the within-county allocation rules into within-market variation in effective entry capacity. The exclusion restriction in the baseline specification is therefore that county area and the three policy-interacted rank components affect incumbent outcomes only through realized competitor entry rather than through omitted local demand shocks.

3.3 Identification, Estimand, and Threats to Validity

The fixed effects absorb time-invariant differences across CDPs and common time shocks. Standard errors are clustered at the CDP level. In a fixed-effects panel IV design, identification comes from variation that remains after those time-invariant market differences are removed (Wooldridge, 2010; Angrist and Pischke, 2009). Because the instruments shift realized entry through licensing rules and cap changes, the resulting 2SLS estimates should be interpreted as local average treatment effects for jurisdictions whose realized entry responds to those policy-driven changes (Imbens and Angrist, 1994). The relevant compliers are therefore markets in which administrative changes in license availability translate into realized entry, while the monotonicity assumption is that more permissive licensing weakly increases the number of active establishments.

The identifying variation in the baseline specification comes from two related but asymmetrically used sources. The cross-county component is the static county-area term. The within-county allocation measures enter as the policy-interacted rank components $Rank_j D_{jt}$, $\mathbb{1}\{Rank_j = 1\} D_{jt}$, and $\mathbb{1}\{Rank_j = 2\} D_{jt}$. That distinction matters because once CDP fixed effects are included, identification must come from within-market changes over time rather than from static cross-sectional differences alone (Wooldridge, 2010; Arellano and Bover, 1995). In this setting, the post-2016 interactions convert the underlying within-county allocation rules into within-market variation in effective entry capacity. Put differently, the policy change is not just an additional source of power; it is the channel that makes the rank-based allocation measures informative in the fixed-effects panel design.

This interpretation matters for both estimands and threats to validity. The coefficient on entry should be read as a LATE for markets whose realized entry responds to administrative changes in license availability (Imbens and Angrist, 1994; Angrist and Pischke, 2009). It is not automatically an average treatment effect for all jurisdictions or for all potential entrants. That framing also clarifies why the most important threats to validity are not generic omitted variables, but specific failures of exclusion: policy changes could be correlated with broader demand growth, and entry near binding caps could reflect rationing or latent profitability rather than competition alone.

The empirical strategy addresses those concerns in three ways. First, the preferred specification restricts the excluded instruments to the log of the square root of county area and the three policy-interacted rank components that map directly into the allocation and cap-expansion rules. Second, we compare specifications that isolate the post-2016 interacted variation from specifications that combine it with static allocation variables. That comparison is useful precisely because the fixed-effects design should reward instruments that deliver within-market movement and downweight instruments that only proxy for cross-market differences. Third, we report pre-expansion balance checks for the initial-allocation components. These steps mitigate most of the concerns and tie the main results more directly to the quasi-mechanical policy variation.

4 Results

This section begins with the main estimates for variable profits and revenue, then turns to the quantity and price margins that underlie those outcomes. We next compare the IV estimates to standard fixed-effects regressions without instruments in order to illustrate the direction of bias. The section then turns to heterogeneity, the role of fixed effects, the first-stage evidence, and spatial robustness to alternative market definitions.

4.1 Impact of Entry on Variable Profit, Revenue, and Their Components

Table 4 reports the 2SLS estimates for the effect of establishment entry on monthly variable profits and revenue using equation (1). The number of establishments in a CDP is instrumented in each specification. In columns 1 and 2, the dependent variable is the natural log of variable profit. In columns 3 and 4, the dependent variable is the natural log of total revenue. Columns 2 and 4 add the Oregon-border indicator, and the coefficient on entry is stable across specifications.

The estimates imply that one additional establishment in the market reduces the monthly variable profits of an existing establishment by approximately 4.1 to 4.2% on average. Both estimates are statistically significant at the 1% level. The same entrant reduces monthly revenue by approximately 3.5 to 3.6% on average, with both estimates statistically significant at the 5% level. The positive and significant coefficient on population within one mile indicates that immediate local demand conditions remain an important determinant of store performance.

These estimates show that additional entry lowers incumbent profits in this setting, but the magnitudes are modest relative to what one might expect in a less regulated retail market. That pattern is consistent with the institutional features of Washington's cannabis market. Retailers continue to possess substantial market power (Hollenbeck and Uetake, 2021; Thomas, 2019), and even after the 2016 cap expansion the modal license allocation across CDPs remains small. In that environment, entry matters, but the effect of each additional entrant can be dampened by regulation and by underlying demand growth.

Table 5 decomposes the revenue effect into quantity and price responses. Column 1 uses the natural log of quantity sold for smokeable flower products as the dependent variable. Column 2 uses the natural log of quantity sold for all other products sold by the unit. Columns 3 and 4 use average price measures. One additional entrant reduces an incumbent's monthly quantity sold of flower by about 3.4% and reduces quantity sold of non-flower products by about 3.7%. Those coefficients are statistically significant at the 5% and 1% levels, respectively. By contrast, the price coefficients are close to zero and not statistically significant.

The decline in variable profit and revenue therefore appears to operate through quantity rather than through price. Evidence on price responses to entry is mixed. Busso and Galiani (2019) find that entry lowers prices, whereas Bergquist and Dinerstein (2020) find negligible price effects. Hollenbeck et al. (2024) further show that assortment differentiation can soften price competition. One interpretation is that Washington retailers operate in a differentiated and still highly regulated market in which posted prices are influenced more by wholesale cost changes and product positioning than by local entry alone. That interpretation is also consistent with the broader evidence in Hollenbeck and Uetake (2021) and with the pricing patterns documented in Escudero (2018).

To the extent that external validity is limited, the likely source is the institutional environment rather than the sign of the effect (Roe and Just, 2009). In this setting, the most plausible moderators are market power, licensing restrictions, and local pricing behavior. Related work on this market indicates that learning is slow and incomplete and that several pricing behaviors limit how quickly retailers adjust. In particular, asymmetric pricing responses, simple pricing heuristics, and independent pricing behavior can keep stores from fully incorporating local demand information, which in turn contributes to losses in variable profits over time (Author blinded for review, 2026). Interviews with retailers are also consistent with the view that managers pay closer attention to wholesale prices and perceived product quality than to local entry counts when setting prices. Those features of the market help explain why quantity adjustments appear first and price responses are muted.

4.2 Treatment Heterogeneity by Regulatory Slack and Baseline Competition

The average estimates above pool together markets that faced materially different licensing constraints. In Washington, the economically relevant margin is not simply that the state revised the aggregate retail cap in 2016, but whether a store's licensing jurisdiction was already close to its allotted number of licenses. We therefore construct a jurisdiction-level cap measure that combines the original allotments, the March 2016 revised allotments, and the relevant city versus county-at-large licensing jurisdiction for each observed market. In the lottery sample, roughly one-third of jurisdiction-months are exactly at the applicable cap and roughly two-thirds are within one license of it. Binding-cap conditions are therefore common rather than exceptional.

Our first heterogeneity specification augments equation (1) with an interaction between realized entry and an indicator for binding-cap jurisdictions:

$$\ln(y_{rt}) = \beta_1 N_{pt} + \beta_2 \left(N_{pt} \times BindingCap_p\right) + \gamma X_r + \alpha_p + \xi_m + \lambda_y + \varepsilon_{rt}. \tag{4}$$

Here, $BindingCap_p$ equals one for jurisdictions that were at or within one license of their relevant cap by the end of 2015. The coefficient β_1 gives the entry effect in higher-slack jurisdictions, while $\beta_1 + \beta_2$ gives the corresponding effect in binding-cap jurisdictions. Because the interaction enters through the endogenous entry variable, we instrument the interacted term with the corresponding interactions between the excluded instruments and the binding-cap indicator. As in the baseline specification, these regressions retain the organizational-size control so that the heterogeneity estimates are not driven mechanically by differences between single-store and multi-store firms.

Table 6 reports the main heterogeneity results. The estimates show that the effect of entry on variable profit is substantially larger in binding-cap jurisdictions than in higher-slack jurisdictions. In the higher-slack group, one additional entrant lowers incumbent variable profit by about 2.8%. In jurisdictions that were already near the cap before the 2016 expansion, the implied effect is a decline of approximately 7.5%. Column 2 shows the same pattern for revenue: the effect is about a 2.4% decline in higher-slack jurisdictions and about a 5.6% decline in binding-cap jurisdictions. This pattern is consistent with the institutional setting. Where licensing constraints were already tight, the entrants induced by the policy change represented a larger relaxation of local scarcity, so each additional entrant had a larger effect on incumbent performance.

Table 7 shows that the same heterogeneity appears in the underlying margins. The quantity response is notably larger in binding-cap jurisdictions, while the price response remains small in both groups. In higher-slack jurisdictions, one additional entrant reduces flower quantity sold by about 2.3% and unit quantity sold by about 2.7%. In binding-cap jurisdictions, the implied quantity effects are larger, with declines of approximately 5.1% for flower and 5.9% for unit sales. By contrast, the price effects remain close to zero in both groups. The implication is that the main heterogeneity operates through quantity rather than price, with the strongest quantity response arising where regulation had made local entry especially scarce.

We treat baseline competition as a complementary heterogeneity exercise because the cap-based measure maps more directly into the policy mechanism. The corresponding specification is

$$\ln(y_{rt}) = \theta_1 N_{pt} + \theta_2 \left(N_{pt} \times BaselineComp_p\right) + \gamma X_r + \alpha_p + \xi_m + \lambda_y + \upsilon_{rt}, \tag{5}$$

where $BaselineComp_p$ is the number of additional firms present in the market at the end of 2015. In this parameterization, θ_1 is the effect of entry in a market with zero baseline additional competitors, and $\theta_1 + \theta_2 b$ gives the effect at baseline competition level b. As in the cap-based specification, the interacted endogenous term is instrumented with the corresponding interactions between the excluded instruments and the predetermined heterogeneity measure.

Table 8 uses the number of additional firms in the market at the end of 2015 as a pre-period measure of baseline competition. The estimated entry effect is largest in thin markets and attenuates modestly as pre-existing competition rises. At zero baseline competitors, the estimated effect is about -10.4% for variable profit and -9.7% for revenue. Evaluated at the mean baseline competition level in the lottery sample, the implied effects are about -8.5% for variable profit and -7.9% for revenue. We interpret this pattern more cautiously than the cap-based results, but it points in the same direction: new entry matters most where local competition was initially weakest.

4.3 Robustness to Alternative Market Definitions

Because the baseline market definition uses CDPs and closely related licensing jurisdictions, a natural concern is that some stores may sit near administrative boundaries even though their effective competitive environment extends across them. To address that concern, we conduct a boundary audit using the geocoded retail locations in the ArcGIS project. Each retail point is first assigned to its containing CDP and county. We then project the retail, place, and county layers into a common feet-based coordinate system, convert the place and county polygons into boundary lines, and measure each store's planar distance to the nearest CDP and county boundary.

The audit shows that boundary exposure is concentrated in CDP assignment rather than county assignment. Of the 566 retail locations in the geocoded audit file, 143 are within 500 feet of a CDP boundary and 207 are within 1,000 feet. By contrast, no stores are within 500 feet of a county boundary and only 2 are within 1,000 feet. These counts suggest that county-level assignment is not an important source of measurement error in this setting, while the CDP results indicate that a nontrivial minority of stores operate close enough to a jurisdiction edge to justify a targeted robustness exercise.

We therefore report two boundary-based robustness checks. The first excludes stores within 500 feet, and then within 1,000 feet, of a CDP boundary from the estimation sample. The second keeps all store-month outcomes in the sample but recomputes the endogenous competition measure so that boundary-adjacent stores no longer count toward the CDP-month entrant count. Conceptually, the latter is the sharper test when the concern is that boundary stores are imperfectly assigned to within-CDP competition rather than that their own outcomes are noisy. In notation, this robustness replaces the baseline competitor count N_{pt} in equation (1) with N_{pt}^{NB}, where only non-boundary stores contribute to measured CDP-month competition. Table 12 and Table 13 report the corresponding estimates.

Both exercises point to the same conclusion. When boundary-adjacent stores are excluded from the estimation sample, the coefficient on entry remains close to the baseline estimate for variable profit, revenue, quantity sold, and price. For example, the variable-profit coefficient moves from about -0.041 in the baseline sample to -0.045 in the 500-foot exclusion and -0.048 in the 1,000-foot exclusion, while the quantity coefficient remains in the neighborhood of -0.034 to -0.040 and the price effect remains close to zero. When the competition variable is instead recomputed so that boundary stores no longer contribute to the CDP competition count, the profit and revenue coefficients remain essentially unchanged, and the quantity and price results do as well. These checks suggest that the main findings are not an artifact of how boundary-adjacent stores are assigned to local competition measures.

As a broader spatial check, we also replace the CDP-based competition measure with drive-time competitor counts. This alternative follows the accessibility literature's emphasis on travel-cost-based market reach and cumulative-opportunity measures (Hansen, 1959; Geurs and van Wee, 2004). It is also consistent with evidence that cannabis demand can cross administrative borders when nearby legal outlets become accessible and with related work showing that consumer purchases in regulated vice markets can respond to travel costs (Hansen et al., 2020; Hobday et al., 2017). Using the ArcGIS Network Analyst origin-destination cost matrix, we count active competing retailers within 10, 15, and 20 minutes of each store in each month, so the baseline N_{pt} in equation (1) is replaced by a retailer-specific travel-time competition count N_{rt}^{τ} for threshold $\tau \in \{10, 15, 20\}$. The 15-minute specification serves as the main alternative market definition because it balances overly narrow and overly broad competitor sets in this geography, while the 10- and 20-minute specifications are reported as sensitivity checks around that central threshold rather than as separate preferred designs. Figure 3 provides the statewide geographic context for those alternative market definitions. Figure 4 then compares the baseline CDP market to the 15-minute drive-time market for the same focal retailer, and Figure 5 shows how the 10-, 15-, and 20-minute thresholds expand that retailer's local competitor set. Those figures illustrate the geographic membership implied by each market definition for a fixed cross-section, while the regressions themselves use month-specific counts of active competing stores within those travel-time bands. Table 14 reports the corresponding estimates.

Under conventional 2SLS, the drive-time specifications preserve the paper's main qualitative results. Across the 10-, 15-, and 20-minute definitions, an increase in nearby competition reduces variable profit and quantity sold, while the estimated effect on price per gram remains small and imprecise. In the 15-minute specification, the coefficient on the drive-time competitor count is approximately -0.76 for log variable profit and -0.61 for log quantity sold, with the corresponding 2SLS price coefficient close to zero. The 10-minute and 20-minute variants show the same directional pattern. Because the drive-time variables count competitors in a different market geometry than the CDP measure, these coefficients are not directly comparable in magnitude to the baseline CDP coefficients. The more informative comparison is the persistence of economically meaningful effects on profits and quantities under alternative spatial definitions of local competition.

At the same time, the first stage is weaker under the drive-time market definitions than in the baseline CDP specifications. Across the reported regressions, the Kleibergen-Paap Wald F statistics generally fall in the 5 to 7 range. For that reason, we also examine weak-IV-robust Anderson-Rubin confidence intervals for the main 15-minute specification. Those intervals remain strictly negative for variable profit and quantity sold, which strengthens the conclusion that the drive-time profitability and quantity results are not an artifact of weak first-stage concerns. The corresponding price interval is strictly positive and excludes zero, in contrast to the conventional near-zero 2SLS estimate. We therefore view the drive-time price regression as especially suggestive and continue to rely on the stronger baseline CDP design for conclusions about the pricing margin. Substantively, however, the broader message is the same as in the boundary-based checks: the paper's central conclusions about profits and quantities do not depend on a knife-edge use of CDP borders, even though the administratively defined CDP framework remains the preferred specification because it aligns most closely with the institutional entry restrictions.

5 Conclusion

This paper provides quasi-experimental IV evidence on how local retail entry affects incumbent variable profits in Washington's recreational cannabis market. Using policy-driven variation from license allocation rules, lottery-based assignment, and the 2016 cap expansion, we estimate that additional entry reduces both variable profits and revenue for incumbent stores. The decline is driven by lower quantities sold rather than lower prices.

The paper adds to the literature by bringing direct profit measures to the study of entry in a regulated retail setting. Existing studies often observe prices or revenue, but not the wholesale purchases needed to construct variable profits. The Washington seed-to-sale data make that possible and show that the quantity margin is central in this market. The heterogeneity results further show that the effect of entry is not uniform across markets. Entry has its largest effect in jurisdictions that were already close to their relevant license caps before the 2016 expansion, and that heterogeneity operates primarily through quantity rather than price.

The baseline estimates survive boundary-based sample exclusions, competition-count redefinitions that exclude boundary-adjacent stores, and alternative drive-time competitor counts built from 10-, 15-, and 20-minute travel-time bands. The drive-time specifications are not strong enough to replace the baseline design, but they point in the same qualitative direction: additional nearby entry reduces incumbent profits and quantities without producing a comparably precise price response.

In many retail settings, the leading margin is price. Here, the data instead indicate that entry reallocates sales and variable profits across stores while prices remain comparatively sticky. We view that difference as an external-validity point rather than as a contradiction. The internal evidence for this setting is strong, but the policy relevance of the magnitude depends on whether other markets share the same combination of licensing restrictions, residual market power, product differentiation, and pricing institutions (Roe and Just, 2009). In other words, the paper identifies a credible effect of entry on variable profits, but it also shows that entry need not operate through price in every market where policy expands competition.

The main limitations point in four directions. First, although the identifying variation is useful, the most credible within-market component still comes from a relatively small number of institutional changes over the sample period. Second, retailer exit is limited over the observed horizon. That feature narrows the scope of the estimates rather than undermining their internal validity: the results speak most directly to how incumbent stores absorb additional competition along intensive margins while remaining in the market, rather than to a long-run equilibrium in which entry and exit jointly determine market structure. Third, the heterogeneity analysis centers on baseline competition and regulatory slack rather than finer spatial exposure to nearby entrants. Fourth, external validity is necessarily tied to a market with explicit license restrictions and substantial residual market power. Even so, this paper distinguishes between average entry effects and the larger effects that arise when policy changes relax binding local entry constraints.

Appendix

A Comparison of OLS and IV Estimates

Table 9 compares the instrumental-variables estimates to ordinary least squares estimates with fixed effects. Columns 1 and 2 use the natural log of variable profit as the dependent variable. Column 1 reports OLS estimates, while column 2 reports the IV estimates from Table 4. The OLS coefficient is statistically significant but smaller in magnitude than the IV coefficient, which suggests upward bias in OLS estimates, attenuating the true estimates. That pattern is consistent with entry responding positively to unobserved local demand conditions that also raise incumbent profits.

Columns 3 and 4 repeat the comparison for the natural log of total revenue.

B Importance of Fixed Effects in Estimation

Equation (1) uses a multi-way fixed-effects specification, and Table 10 shows why those controls matter. Without fixed effects (column 1), the estimated relationship is large and significant. Adding year fixed effects (column 2) sharply attenuates the estimate, which points to important common shocks during a period of rapid market development and regulatory change. Using CDP fixed effects without time effects (column 3) instead strengthens the estimate, indicating that the identifying comparisons come from within-market changes over time rather than cross-sectional differences across markets.

The comparison also clarifies why the preferred instrument set must contain within-market variation. Once CDP fixed effects absorb time-invariant market differences, static allocation variables alone have little identifying power. In split-source checks, the post-2016 interacted rank instruments remain close to the combined baseline, whereas specifications relying only on initial-allocation components become weak. This is standard panel-IV logic: with local fixed effects, identification must come from policy-driven within-market changes in effective license availability (Wooldridge, 2010; Angrist and Pischke, 2009). Consistent with that interpretation, Table 11 reports a preferred Kleibergen-Paap Wald F-statistic above 20, and pre-2016 balance checks show little relationship between the initial-allocation components and pre-expansion outcomes.

References

Jean Marie Abraham, Martin Gaynor, and William B. Vogt. Entry and competition in local hospital markets. *The Journal of Industrial Economics*, 55(2):265–288, 2007. doi: 10.1111/j.1467-6451.2007.00311.x.

Joshua D. Angrist and Alan B. Krueger. Instrumental variables and the search for identification: From supply and demand to natural experiments. *Journal of Economic Perspectives*, 15(4):69–85, 2001. doi: 10.1257/jep.15.4.69.

Joshua D Angrist and Jörn-Steffen Pischke. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press, 2009.

Manuel Arellano and Olympia Bover. Another look at the instrumental variable estimation of error-components models. *Journal of Econometrics*, 68(1):29–51, 1995.

Author blinded for review. Blinded reference omitted for peer review. Manuscript available upon request, 2026.

Lauren Falcao Bergquist and Michael Dinerstein. Competition and entry in agricultural markets: Experimental evidence from kenya. *American Economic Review*, 110(12):3705–47, 2020.

Steven T. Berry. Estimation of a model of entry in the airline industry. *Econometrica*, 60(4):889–917, 1992. doi: 10.2307/2951571.

Steven T. Berry and Joel Waldfogel. Free entry and social inefficiency in radio broadcasting. *The RAND Journal of Economics*, 30(3):397–420, 1999. doi: 10.2307/2556055.

Timothy F Bresnahan and Peter C Reiss. Entry and competition in concentrated markets. *Journal of political economy*, 99(5):977–1009, 1991.

Matias Busso and Sebastian Galiani. The causal effect of competition on prices and quality: Evidence from a field experiment. *American Economic Journal: Applied Economics*, 11(1):33–56, 2019.

Jonathan P Caulkins and Linden Dahlkemper. Retail store allocation. Technical report, BOTEC Analysis Corporation for the Washington State Liquor Control Board, jun 2013. URL https://lcb.wa.gov/publications/Marijuana/BOTECreports/Re_Store_Allocation_Task_Report-Final.pdf.

James M. Cole. Guidance regarding marijuana enforcement. https://www.justice.gov/iso/opa/resources/3052013829132756857467.pdf, 2013. Memorandum, U.S. Department of Justice.

Matias Escudero. Rule-of-thumb pricing: Retail cannabis in washington state. Technical report, Mimeo, Northwestern University, 2018.

Karst T. Geurs and Bert van Wee. Accessibility evaluation of land-use and transport strategies: Review and research directions. *Journal of Transport Geography*, 12(2):127–140, 2004.

Benjamin Hansen, Keaton Miller, and Caroline Weber. Getting into the weeds of tax invariance. Technical report, National Bureau of Economic Research, 2017.

Benjamin Hansen, Keaton Miller, and Caroline Weber. Federalism, partial prohibition, and cross-border sales: evidence from recreational marijuana. *Journal of Public Economics*, 187:104159, 2020.

Benjamin Hansen, Keaton Miller, and Caroline Weber. Vertical integration and production inefficiency in the presence of a gross receipts tax. *Journal of Public Economics*, 212:104693, 2022. doi: 10.1016/j.jpubeco.2022.104693.

Walter G. Hansen. How accessibility shapes land use. *Journal of the American Institute of Planners*, 25(2): 73–76, 1959.

Michelle Hobday, E Lensvelt, E Gordon, W Liang, L Meuleners, and T Chikritzhs. Distance travelled to purchase alcohol and the mediating effect of price. *Public health*, 144:48–56, 2017.

Brett Hollenbeck and Renato Zaterka Giroldo. Winning big: Scale and success in retail entrepreneurship. *Marketing Science*, 41(2):271–293, 2022.

Brett Hollenbeck and Kosuke Uetake. Taxation and market power in the legal marijuana industry. *The RAND Journal of Economics*, 52(3):559–595, 2021.

Brett Hollenbeck, Sylvia Hristakeva, and Kosuke Uetake. Retailer price competition and assortment differentiation: Evidence from entry lotteries. Technical report, Working Paper, 2024.

Guido W Imbens and Joshua D Angrist. Identification and estimation of local average treatment effects. *Econometrica*, 62(2):467–475, 1994.

Steven Manson, Jonathan Schroeder, David Van Riper, Tracy Kugler, and Steven Ruggles. *IPUMS National Historical Geographic Information System: Version 16.0 [dataset]*. IPUMS, Minneapolis, MN, 2021.

Brian E. Roe and David R. Just. Internal and external validity in economics research: Tradeoffs between experiments, field experiments, natural experiments, and field data. *American Journal of Agricultural Economics*, 91(5):1266–1271, 2009.

State of Washington House of Representatives. Initiative 502. https://leg.wa.gov/CodeReviser/documents/INITIATIVE502.SL.pdf, 2012. Accessed: February 10, 2022.

State of Washington Liquor Control Board. Press Release: BioTrackTHC. https://lcb.wa.gov/pressreleases/liquor-control-board-announces-biotrackthc-provide-traceability-software, 2013a. Accessed: February 10, 2022.

State of Washington Liquor Control Board. Press releases. https://lcb.wa.gov/pressreleases/2013-press-releases, 2013b. Accessed: February 10, 2022.

Betsey Stevenson and Justin Wolfers. *Principles of Economics*. Worth Publishers, 2020.

Danna Thomas. License quotas and the inefficient regulation of sin goods: Evidence from the washington recreational marijuana market. SSRN working paper 3312960, 2019.

Washington State Liquor and Cannabis Board. Liquor control board approves lottery process for retail cannabis stores. https://lcb.wa.gov/pressreleases/liquor-control-board-approves-lottery-process-retail-cannabis-stores, 2014. Accessed: February 10, 2022.

Washington State Liquor and Cannabis Board. LCB to increase number of retail marijuana stores. https://lcb.wa.gov/pressreleases/lcb-to-increase-number-of-retail-mj-stores, 2015. Accessed: February 10, 2022.

Washington State Liquor and Cannabis Board. Proposed number of retail stores to serve the medical marijuana market. https://lcb.wa.gov/sites/default/files/publications/Cannabis/MJ_Retail_Allocation_3-8-16.pdf, 2016. Revised store allocation, March 8, 2016.

Jeffrey M Wooldridge. *Econometric analysis of cross section and panel data*. MIT press, 2010.

C Tables

C.1 Table 1: Establishment Entry Over Time

Year	Beginning of Year	End of Year
2014	17*	86
2015	98	199
2016	202	337
2017	345	381**

Notes: The beginning-of-year and end-of-year counts are based on locations reporting sales in the traceability data.
* The first observations are recorded in July 2014.
** The dataset ends before December 2017.

C.2 Table 2: Organizations with Multiple Establishments

Year	Organizations with 2+ Stores	Maximum Stores per Organization
2014	6	2
2015	11	2
2016	54	3 (two firms)
2017	82	3 (six firms)

Notes: Column (2) reports the number of organizations operating at least two retail establishments in a given year. Column (3) reports the largest observed number of establishments linked to a single organization ID, with the number of organizations attaining that maximum shown in parentheses.

C.3 Table 3: License Allotments and CDP Population Ranks

Jurisdiction	Retail Outlets	County Population Rank	2010 County Population
Mill Creek	1	18	713,335
Ellensburg	2	1	40,915
Bainbridge Island	1*	2	251,133
University Place	1	5	795,225
Mount Vernon	3	1	116,901
Walla Walla	2	1	58,781
Bellingham	6	1	201,140
Spokane Valley	3	2	471,221

Notes: Retail outlets indicate the number of licenses allotted.
* The rank-1 city in this county received two licenses.

C.4 Table 4: Entry Effects on Variable Profit and Revenue

	Variable Profit (1)	Variable Profit (2)	Revenue (3)	Revenue (4)
Number of Establishments- CDP	-0.042***	-0.041***	-0.036**	-0.036**
	(0.015)	(0.014)	(0.016)	(0.016)
Log Population- 1mi Buffer	0.353***	0.352***	0.102	0.102
	(0.127)	(0.128)	(0.182)	(0.183)
Log Population- 2mi Buffer	-0.211	-0.211	0.099	0.099
	(0.191)	(0.192)	(0.272)	(0.272)
Number of Locs- Organization	-0.153	-0.153	-0.245	-0.244
	(0.203)	(0.204)	(0.207)	(0.207)
Border Stores Open- Post-2015m9		-0.101		-0.014
		(0.179)		(0.293)
Observations	7,357	7,357	7,461	7,461

Notes: Each column reports a retailer-month 2SLS regression. The endogenous regressor is the number of establishments reporting sales in the retailer's CDP-month. The dependent variable is log variable profit in columns (1) and (2) and log revenue in columns (3) and (4). All specifications include CDP, month-of-year, and year fixed effects. Controls include population within one mile, population within two miles, and the number of establishments in the retailer's organization; columns (2) and (4) also include the Oregon-border post-legalization indicator. The corresponding Kleibergen-Paap Wald F-statistics are 13.59, 20.56, 15.08, and 21.54 for columns (1) through (4), respectively. Standard errors are clustered at the CDP level and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C.5 Table 5: Entry Effects on Quantity Sold and Price

	Quantity - Flower (1)	Quantity - units (2)	Price Per Gram (3)	Price Per Unit (4)
Number of Locations- CDP	-0.034**	-0.037***	-0.001	-0.002
	(0.014)	(0.014)	(0.002)	(0.003)
Log Population- 1mi Buffer	0.356**	0.213	-0.036	0.014
	(0.157)	(0.138)	(0.029)	(0.022)
Log Population- 2mi Buffer	-0.268	-0.047	0.044	-0.018
	(0.214)	(0.196)	(0.034)	(0.034)
# Locations- Organization	-0.089	-0.177	-0.019	-0.027
	(0.180)	(0.194)	(0.015)	(0.020)
Border Stores Open- Post-2015m9	0.179	0.229	-0.063	-0.147*
	(0.141)	(0.151)	(0.075)	(0.082)
Observations	7,451	7,461	7,451	7,461

Notes: Each column reports a retailer-month 2SLS regression. The endogenous regressor is the number of establishments reporting sales in the retailer's CDP-month. Columns (1) and (2) use log quantity sold for flower and unit-based products. Columns (3) and (4) use log average retail price per gram and log average retail price per unit. All specifications include the baseline controls together with CDP, month-of-year, and year fixed effects. The corresponding Kleibergen-Paap Wald F-statistics are 33.78, 21.54, 33.78, and 21.54 for columns (1) through (4), respectively. Standard errors are clustered at the CDP level and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C.6 Table 6: Heterogeneity by Binding Cap and Regulatory Slack

	Variable Profit (1)	Revenue (2)
Entry Effect in Higher-Slack Jurisdictions	-0.028***	-0.024***
	(0.003)	(0.004)
Additional Entry Effect in Binding-Cap Jurisdictions	-0.046*	-0.032
	(0.024)	(0.033)
Log Population- 1mi Buffer	0.360***	0.101
	(0.123)	(0.184)
Log Population- 2mi Buffer	-0.212	0.102
	(0.189)	(0.275)
Number of Locs- Organization	-0.127	-0.230
	(0.212)	(0.217)
Border Stores Open- Post-2015m9	-0.046	0.028
	(0.153)	(0.273)
Observations	7,095	7,195
Implied Entry Effect: Binding-Cap Jurisdictions	-0.075	-0.056
SE: Implied Entry Effect	0.026	0.035

Notes: Each column reports a retailer-month 2SLS heterogeneity regression in which the entry effect varies by regulatory slack. "Higher-slack jurisdictions" are the omitted group. "Binding-cap jurisdictions" are jurisdictions that were at or within one license of their relevant cap by the end of 2015, using the jurisdiction-specific cap measure constructed from the original and revised retail allotments. Because the cap crosswalk does not match nine place records to a licensing-jurisdiction cap entry, these specifications are estimated on the matched-cap sample. The first coefficient gives the entry effect in higher-slack jurisdictions; the implied binding-cap effect, reported in the lower panel, is the sum of the first two coefficients. All specifications include the baseline controls together with CDP, month-of-year, and year fixed effects. The corresponding Kleibergen-Paap Wald F-statistics are 114.88 and 104.20 for columns (1) and (2), respectively. Standard errors are clustered at the CDP level and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C.7 Table 7: Binding-Cap Heterogeneity in Quantity and Price

	Quantity - Flower (1)	Quantity - Units (2)	Price per Gram (3)	Price per Unit (4)
Entry Effect in Higher-Slack Jurisdictions	-0.023***	-0.027***	-0.001	-0.000
	(0.003)	(0.003)	(0.001)	(0.001)
Additional Entry Effect in Binding-Cap Jurisdictions	-0.028	-0.032	-0.002	-0.005
	(0.030)	(0.029)	(0.005)	(0.005)
Log Population- 1mi Buffer	0.355**	0.212	-0.036	0.014
	(0.158)	(0.139)	(0.029)	(0.022)
Log Population- 2mi Buffer	-0.266	-0.043	0.044	-0.017
	(0.216)	(0.198)	(0.034)	(0.034)
Number of Locs- Organization	-0.067	-0.160	-0.020	-0.024
	(0.188)	(0.202)	(0.015)	(0.021)
Border Stores Open- Post-2015m9	0.219	0.267*	-0.062	-0.141*
	(0.143)	(0.150)	(0.072)	(0.077)
Observations	7,185	7,195	7,185	7,195
Implied Entry Effect: Binding-Cap Jurisdictions	-0.051	-0.059	-0.003	-0.006
SE: Implied Entry Effect	0.031	0.030	0.005	0.006

Notes: Each column reports the retailer-month 2SLS binding-cap heterogeneity specification by outcome. "Higher-slack jurisdictions" are the omitted group, and "binding-cap jurisdictions" are defined as jurisdictions that were at or within one license of their relevant cap by the end of 2015. Because the cap crosswalk does not match nine place records to a licensing-jurisdiction cap entry, these specifications are estimated on the matched-cap sample. The first coefficient gives the entry effect in higher-slack jurisdictions; the implied binding-cap effect, reported in the lower panel, is the sum of the first two coefficients. Columns (1) and (2) use log quantity sold for flower and unit-based products, and columns (3) and (4) use log average retail price per gram and log average retail price per unit. All specifications include the baseline controls together with CDP, month-of-year, and year fixed effects. The corresponding Kleibergen-Paap Wald F-statistics are 191.11, 104.33, 190.31, and 104.20 for columns (1) through (4), respectively. Standard errors are clustered at the CDP level and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C.8 Table 8: Complementary Heterogeneity by Baseline Competition

	Variable Profit (1)	Revenue (2)
Entry Effect When Baseline Competition = 0	-0.104**	-0.097*
	(0.043)	(0.058)
Extra Entry Effect per Baseline Additional Firm	0.004*	0.004
	(0.002)	(0.003)
Log Population- 1mi Buffer	0.356***	0.100
	(0.125)	(0.185)
Log Population- 2mi Buffer	-0.203	0.113
	(0.191)	(0.277)
Number of Locs- Organization	-0.153	-0.264
	(0.223)	(0.228)
Border Stores Open- Post-2015m9	-0.051	0.019
	(0.149)	(0.259)
Observations	6,689	6,788
Mean Baseline Additional Firms	4.516	4.516
Implied Entry Effect at Mean Baseline Competition	-0.085	-0.079
SE: Implied Entry Effect	0.033	0.044

Notes: Each column reports a retailer-month 2SLS heterogeneity regression in which the entry effect varies with baseline competition. The dependent variables are the log of variable profit and the log of revenue. "Baseline additional firms" is the number of additional firms present in the market at the end of 2015. The first coefficient gives the entry effect when baseline competition equals zero; the effect at the sample mean baseline competition is reported in the lower panel. All specifications include the baseline controls together with CDP, month-of-year, and year fixed effects. The corresponding Kleibergen-Paap Wald F-statistics are 164.84 and 156.71 for columns (1) and (2), respectively. Standard errors are clustered at the CDP level and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C.9 Table 9: Comparison of OLS and IV Estimates

	Variable Profit (1)	Variable Profit (2)	Revenue (3)	Revenue (4)
Number of Locations- CDP	-0.014***	-0.041***	-0.015**	-0.036**
	(0.005)	(0.014)	(0.006)	(0.016)
Log Population- 1mi Buffer	0.329**	0.352***	0.084	0.102
	(0.130)	(0.128)	(0.184)	(0.183)
Log Population- 2mi Buffer	-0.184	-0.211	0.119	0.099
	(0.193)	(0.192)	(0.273)	(0.272)
Number of Locs- Organization	-0.173	-0.153	-0.265	-0.244
	(0.207)	(0.204)	(0.210)	(0.207)
Border Stores Open- Post-2015m9	-0.020	-0.101	0.045	-0.014
	(0.200)	(0.179)	(0.307)	(0.293)
Constant	8.953***		9.641***	
	(0.975)		(1.268)	
Observations	7,409	7,357	7,513	7,461

Notes: Columns (1) and (3) report fixed-effects OLS estimates, while columns (2) and (4) report the corresponding baseline 2SLS estimates. The endogenous regressor is the number of establishments reporting sales in the retailer's CDP-month. The dependent variable is log variable profit in columns (1) and (2) and log revenue in columns (3) and (4). All specifications include the baseline controls together with CDP, month-of-year, and year fixed effects. For the IV columns, the corresponding Kleibergen-Paap Wald F-statistics are 20.56 in column (2) and 21.54 in column (4). Standard errors are clustered at the CDP level and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C.10 Table 10: Assessment of Fixed Effects

	Variable Profit (1)	Variable Profit (2)	Variable Profit (3)	Variable Profit (4)	Variable Profit (5)
Number of Locations- CDP	-0.032***	-0.017*	-0.078***	-0.043***	-0.041***
	(0.011)	(0.009)	(0.024)	(0.015)	(0.014)
Log Population- 1mi Buffer	0.008	0.086	0.496***	0.391***	0.352***
	(0.175)	(0.165)	(0.136)	(0.126)	(0.128)
Log Population- 2mi Buffer	0.319	0.183	-0.479**	-0.289	-0.211
	(0.196)	(0.184)	(0.221)	(0.197)	(0.192)
Number of Locs- Organization	-0.334*	-0.071	-0.415*	-0.206	-0.153
	(0.175)	(0.142)	(0.241)	(0.207)	(0.204)
Border Stores Open- Post-2015m10	0.049	0.166	-0.490***	-0.289*	-0.101
	(0.248)	(0.226)	(0.175)	(0.174)	(0.179)
Constant	7.121***				
	(0.548)				
Observations	7,357	7,357	7,357	7,357	7,357

Notes: Each column reports the relationship between entry and log variable profit under a different fixed-effects structure. Entry is treated as endogenous and instrumented in every column. The dependent variable for each column is the log of variable profit. Column (1) includes no fixed effects, column (2) includes year fixed effects only, column (3) includes CDP fixed effects only, column (4) includes CDP and year fixed effects, and column (5) adds the full seasonal controls used in the baseline specification. All columns include population within one mile, population within two miles, organizational size, and the Oregon-border post-legalization indicator. For column (1), the robust excluded-instrument first-stage F-statistic is 11.49. For columns (2) through (5), the corresponding Kleibergen-Paap Wald F-statistics are 10.75, 17.69, 23.23, and 20.56, respectively. Standard errors are clustered at the CDP level and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C.11 Table 11: First Stage Results

	Number of Locations- CDP (1)	Number of Locations- CDP (2)	Number of Locations- CDP (3)
log Sqrt County Area	0.964	14.008	1.973*
	(1.753)	(12.704)	(1.075)
Log Population- 1mi Buffer	0.773***	0.802***	0.371*
	(0.280)	(0.265)	(0.209)
Log Population- 2mi Buffer	-0.846*	-0.891*	-0.543
	(0.476)	(0.455)	(0.363)
# Locations- Organization	-0.156	-0.174	-0.263
	(0.500)	(0.498)	(0.395)
Border Stores Open- Post-2015m9	-2.815	-3.524	-6.132*
	(2.688)	(2.945)	(3.435)
$D_{jt}\times$ rank		-0.217*	-0.014
		(0.125)	(0.014)
$D_{jt}\times$ rank-1 indicator			11.060**
			(4.462)
$D_{jt}\times$ rank-2 indicator			4.476***
			(0.833)
Observations	7,409	7,357	7,357

Notes: Each column reports a first-stage regression with the number of establishments reporting sales in the retailer's CDP-month as the dependent variable. Column (1) uses only log square-root county area as the excluded instrument. Column (2) uses county area together with $D_{jt}\times$ rank. Column (3) uses county area together with $D_{jt}\times$ rank, $D_{jt}\times$ rank-1 indicator, and $D_{jt}\times$ rank-2 indicator, where $D_{jt} = NoBan_j \times Post2016_t \times LicenseDouble_j$. All columns include population within one mile, population within two miles, organizational size, the Oregon-border post-legalization indicator, and the baseline fixed-effects structure. The corresponding Kleibergen-Paap Wald F-statistics are 0.30, 1.63, and 20.56 for columns (1) through (3), respectively. Standard errors are clustered at the CDP level and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C.12 Table 12: Boundary-Exclusion Robustness

	Baseline sample	Exclude 500 ft	Exclude 1,000 ft
Panel A. Variable profit			
Entry coefficient	-0.041^{***}	-0.045^{***}	-0.048^{***}
	(0.014)	(0.016)	(0.018)
Observations	7,357	5,756	5,002
Panel B. Revenue			
Entry coefficient	-0.036^{**}	-0.038^{**}	-0.048^{**}
	(0.016)	(0.019)	(0.023)
Observations	7,461	5,835	5,071
Panel C. Quantity sold			
Entry coefficient	-0.034^{**}	-0.034^{***}	-0.040^{***}
	(0.014)	(0.013)	(0.015)
Observations	7,451	5,825	5,064
Panel D. Price per gram			
Entry coefficient	-0.001	-0.001	-0.002
	(0.002)	(0.002)	(0.003)
Observations	7,451	5,825	5,064

Notes: Each entry reports the coefficient on the endogenous entry measure from a retailer-month 2SLS regression with the baseline fixed-effects structure and baseline controls. Columns (1) through (3) re-estimate the baseline specification after excluding stores within 500 feet and 1,000 feet of a CDP boundary from the estimation sample. Panels A through D report, respectively, log variable profit, log revenue, log quantity sold, and log price per gram. The corresponding Kleibergen-Paap Wald F-statistics are 20.56, 83.75, and 124.20 in Panel A; 21.54, 82.19, and 113.87 in Panel B; 33.78, 81.78, and 113.40 in Panel C; and 33.78, 81.78, and 113.40 in Panel D. Standard errors, clustered at the CDP level, are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C.13 Table 13: Boundary-Based Competition Redefinition Robustness

	Baseline count	Exclude 500 ft from count	Exclude 1,000 ft from count
Panel A. Variable profit			
Entry coefficient	-0.041^{***}	-0.041^{***}	-0.041^{***}
	(0.014)	(0.014)	(0.014)
Observations	7,357	7,357	7,357
Panel B. Revenue			
Entry coefficient	-0.036^{**}	-0.036^{**}	-0.036^{**}
	(0.016)	(0.016)	(0.016)
Observations	7,461	7,461	7,461
Panel C. Quantity sold			
Entry coefficient	-0.034^{**}	-0.035^{***}	-0.035^{**}
	(0.014)	(0.013)	(0.014)
Observations	7,451	7,451	7,451
Panel D. Price per gram			
Entry coefficient	-0.001	-0.001	-0.001
	(0.002)	(0.002)	(0.002)
Observations	7,451	7,451	7,451

Notes: Each entry reports the coefficient on the endogenous competition measure from a retailer-month 2SLS regression with the baseline fixed-effects structure and baseline controls. Columns (2) and (3) keep all store-month outcomes in the sample but redefine the CDP-month competition count so that stores within 500 feet and 1,000 feet of a CDP boundary do not count toward measured CDP competition. Panels A through D report, respectively, log variable profit, log revenue, log quantity sold, and log price per gram. The corresponding Kleibergen-Paap Wald F-statistics are 20.56, 21.87, and 21.14 in Panel A; 21.54, 23.10, and 21.34 in Panel B; 33.78, 34.60, and 32.38 in Panel C; and 33.78, 34.60, and 32.38 in Panel D. Standard errors, clustered at the CDP level, are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C.14 Table 14: Drive-Time Competition Robustness

	10-minute market	15-minute market	20-minute market
Panel A. Variable profit			
Entry coefficient	-1.245^{***}	-0.759^{***}	-0.617^{***}
	(0.440)	(0.274)	(0.214)
Observations	7,357	7,357	7,357
Kleibergen-Paap F	5.75	6.03	7.39
Panel B. Quantity sold			
Entry coefficient	-0.985^{***}	-0.610^{***}	-0.475^{***}
	(0.347)	(0.212)	(0.159)
Observations	7,493	7,493	7,493
Kleibergen-Paap F	5.16	5.49	6.99
Panel C. Price per gram			
Entry coefficient	-0.002	-0.004	-0.011
	(0.037)	(0.025)	(0.022)
Observations	7,451	7,451	7,451
Kleibergen-Paap F	5.15	5.49	6.96

Notes: Each panel reports the coefficient on the endogenous drive-time competitor count from a retailer-month 2SLS regression with the baseline fixed-effects structure and baseline controls. Competition is measured as the number of active competing retailers within 10, 15, or 20 minutes of the focal store in a given month, using ArcGIS Network Analyst origin-destination travel times. Panels A through C report log variable profit, log quantity sold, and log price per gram. The Kleibergen-Paap Wald F-statistics are reported for the corresponding first stage. For the 15-minute specification, the 95% Anderson-Rubin confidence intervals are [-2.516, -1.206] for variable profit, [-2.196, -0.282] for quantity sold, and [0.029, 0.138] for price per gram. The intervals are generated by inverting cluster-robust Anderson-Rubin tests in the active replication code. Standard errors, clustered at the CDP level, are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

D Figure Legends

Figure 1. Retail Cannabis Locations and CDPs. Retail cannabis locations are shown relative to Washington Census Designated Place boundaries as of October 2017.

Figure 2. Retail Locations and Census Tracts in WA, OR, and ID. The figure shows the tract geography used to construct local population controls together with the observed retail locations.

Figure 3. Alternative Drive-Time Competitor Sets. The figure shows statewide 10-, 15-, and 20-minute drive-time competitor sets used in the alternative market-definition exercise.

Figure 4. Baseline CDP Market versus 15-Minute Drive-Time Market. The figure compares the administratively defined CDP market to the 15-minute drive-time competitor set for the same focal retailer.

Figure 5. Zoomed Example of Drive-Time Market Membership. The zoomed comparison shows how the local competitor set expands as the travel-time threshold moves from 10 to 15 to 20 minutes.

E Figures

E.1 Figure 1: Retail Cannabis Locations and CDPs

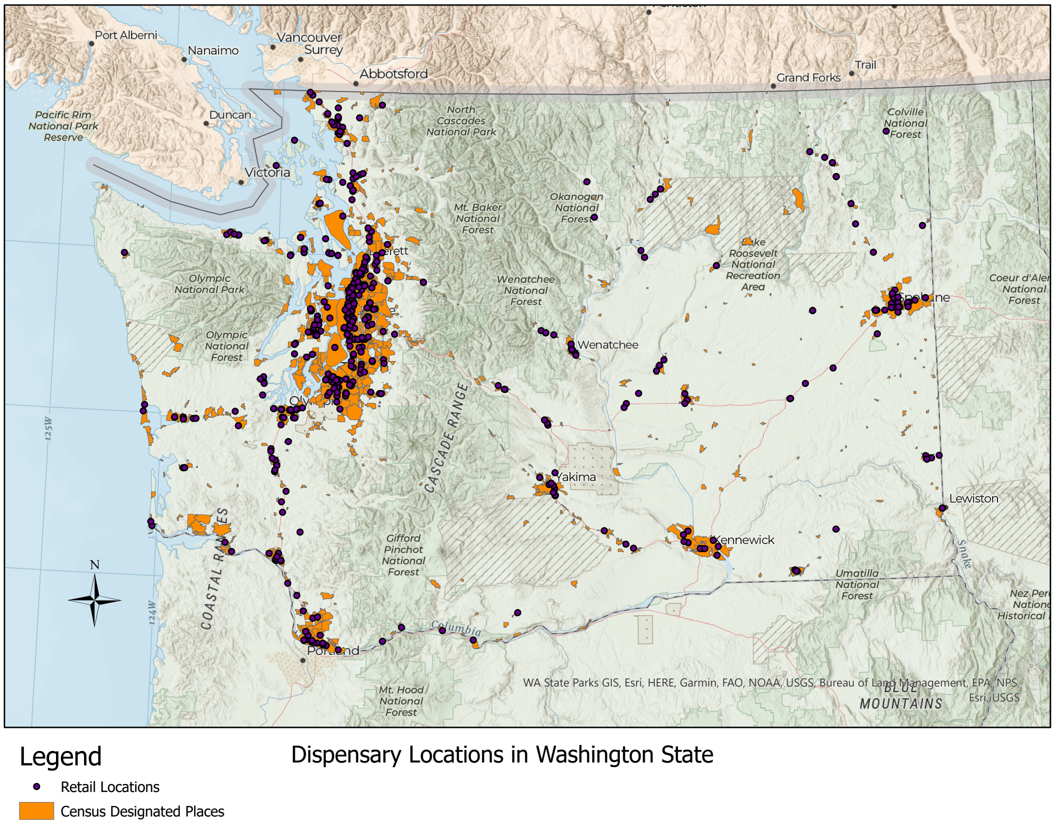

E.2 Figure 2: Retail Locations and Census Tracts in WA, OR, and ID

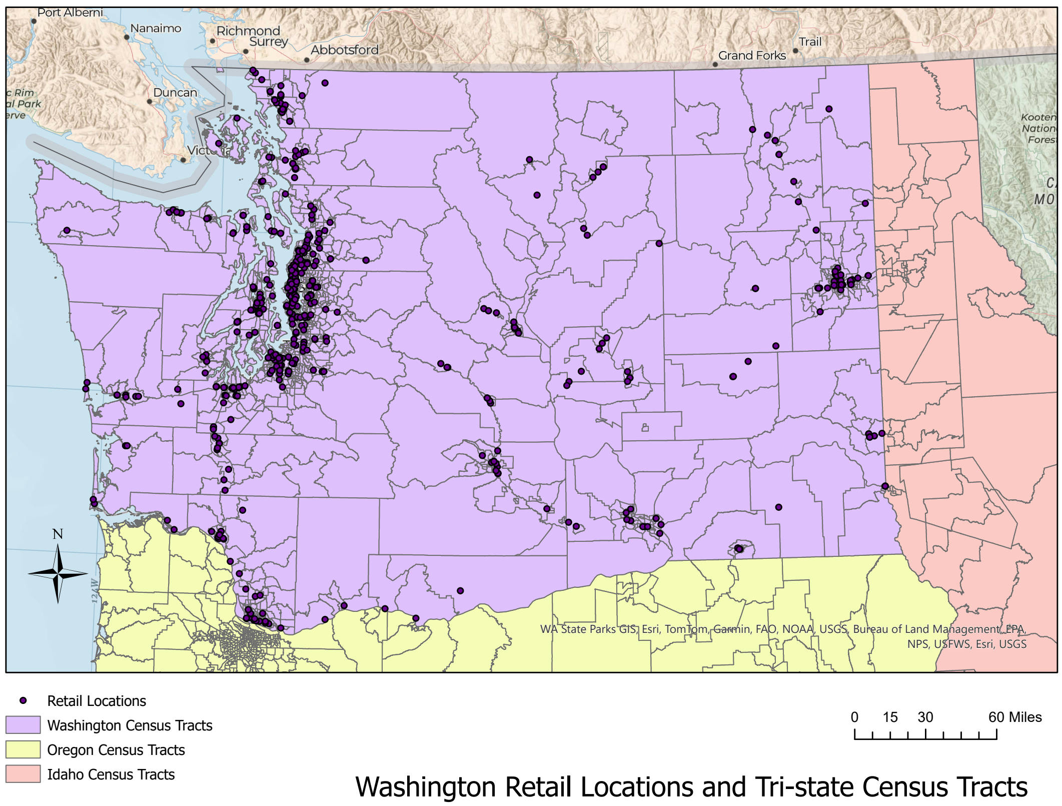

E.3 Figure 3: Alternative Drive-Time Competitor Sets

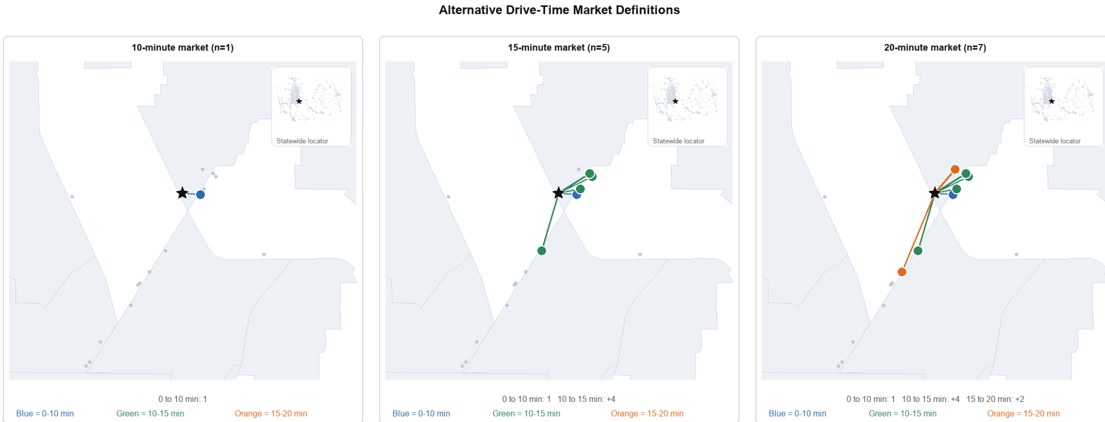

E.4 Figure 4: Baseline CDP Market versus 15-Minute Drive-Time Market

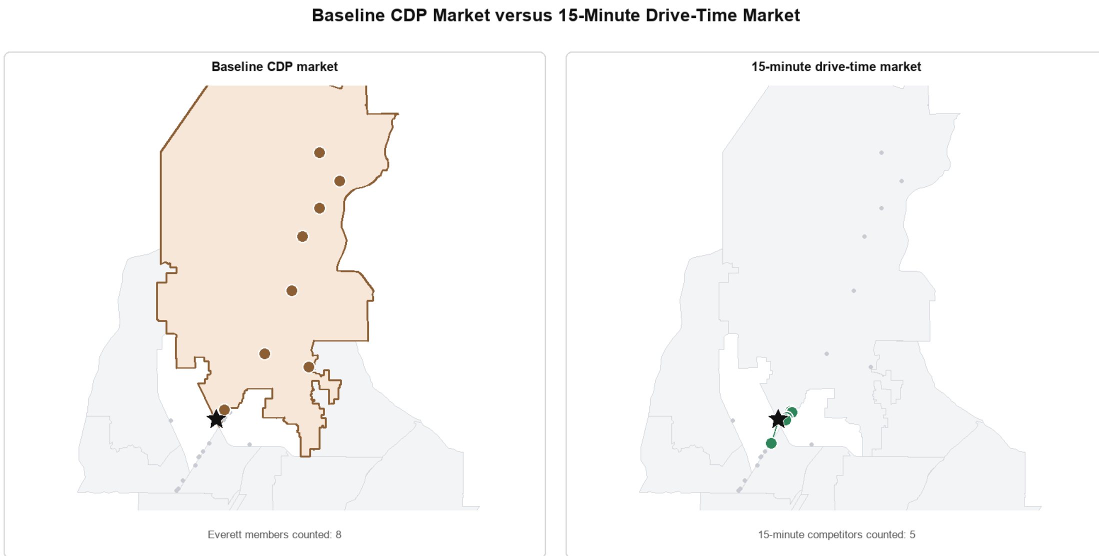

E.5 Figure 5: Zoomed Example of Drive-Time Market Membership