

Supplemental Appendix - The Pass-through of Retail Crime

Carl Hase and Johannes Kasinger*

A Cannabis industry background

A.1 Cannabis consumption and supply in Washington state

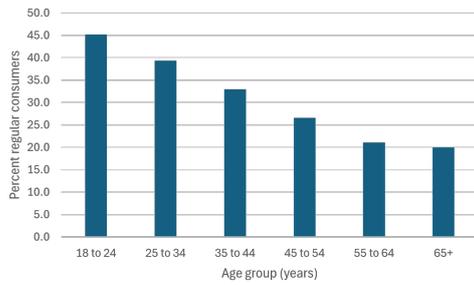
Cannabis consumers Cannabis use is widespread in Washington state. Approximately 30% of Washington adults consume cannabis on a monthly basis (Washington State Department of Health, 2024). For context, about 10% of adults in Washington consume cigarettes. Figure A.1 shows cannabis consumption in Washington state along various demographic lines. The data come from the Behavioral Risk Factor Surveillance System (BRFSS) survey, an annual survey conducted by the Department of Health in partnership with the Centers for Disease Control and Prevention (CDC). The purpose of the survey is to measure changes in the health behaviors of people in Washington state (Washington State Department of Health, 2024).

The figures show the percent of each demographic group that consumes cannabis at least once a month. Panel A illustrates that over 40% of adults age 18 to 24 use cannabis regularly; the same holds for adults age 25 to 34. 33% of adults age 35 to 44 consume regularly, while only one in five adults aged 55+ consume regularly. Panel B shows that regular cannabis consumption decreases monotonically with household income. There are no major differences between levels of education (Panel C). Panel D shows that consumption is highest among American indian, black, multiracial, and other (approx 40%), while consumption is lowest for asian and hispanic adults (approx. 25%). Panel E shows that males consume more than females (35% vs 25%).

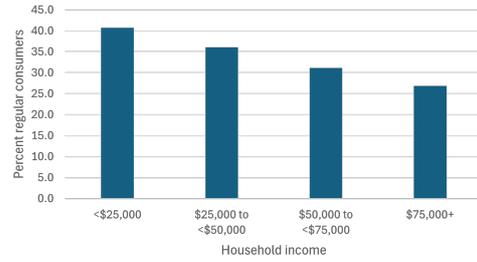
The cannabis supply chain Figure A.2 shows the stages of the cannabis supply chain. Most cannabis in Washington is grown in indoor facilities ranging in size from 2,000 to 30,000 square feet of plant canopy. When plants reach a mature stage, their buds are harvested, dried, and cured. The majority of cannabis is consumed in this unprocessed form (called

*Hase: Goethe University Frankfurt and Johannes Gutenberg University Mainz (carlhase@gmail.com); Kasinger: Tilburg School of Economics and Management - Tilburg University (j.kasinger@tilburguniversity.edu).

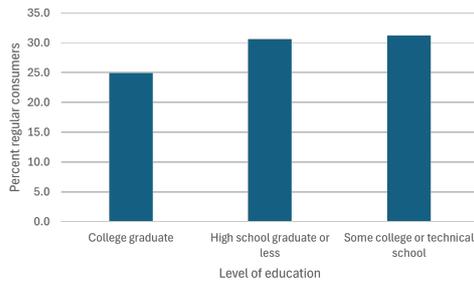
Figure A.1: Demographic characteristics of regular cannabis consumers



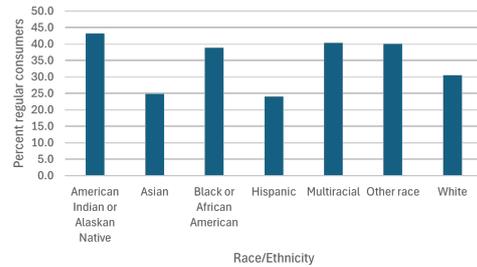
(A) Age



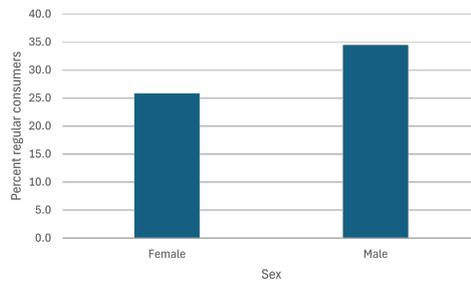
(B) Household income



(C) Education



(D) Race/Ethnicity

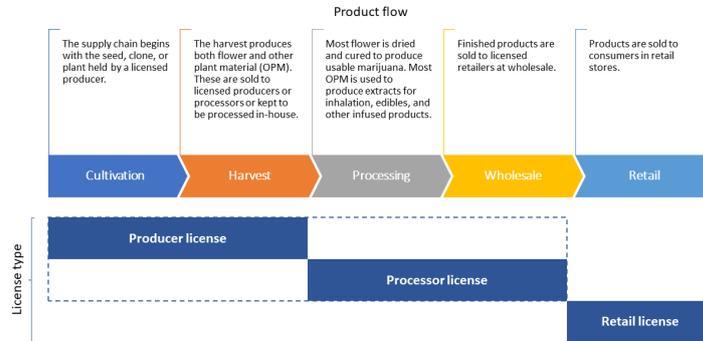


(E) Sex

Notes: This figure presents the distribution of regular cannabis consumers in Washington state, broken down by various demographic characteristics: (A) age, (B) household income, (C) education level, (D) race/ethnicity, and (E) gender. Each bar represents the proportion of regular cannabis consumers within the respective subgroup. The data is from the 2021 Behavioral Risk Factor Surveillance System by the Washington State Department of Health, Center for Health Statistics.

”Usable marijuana”) while the rest is processed into derivative subproducts like edibles and concentrates.

Figure A.2: The cannabis supply chain

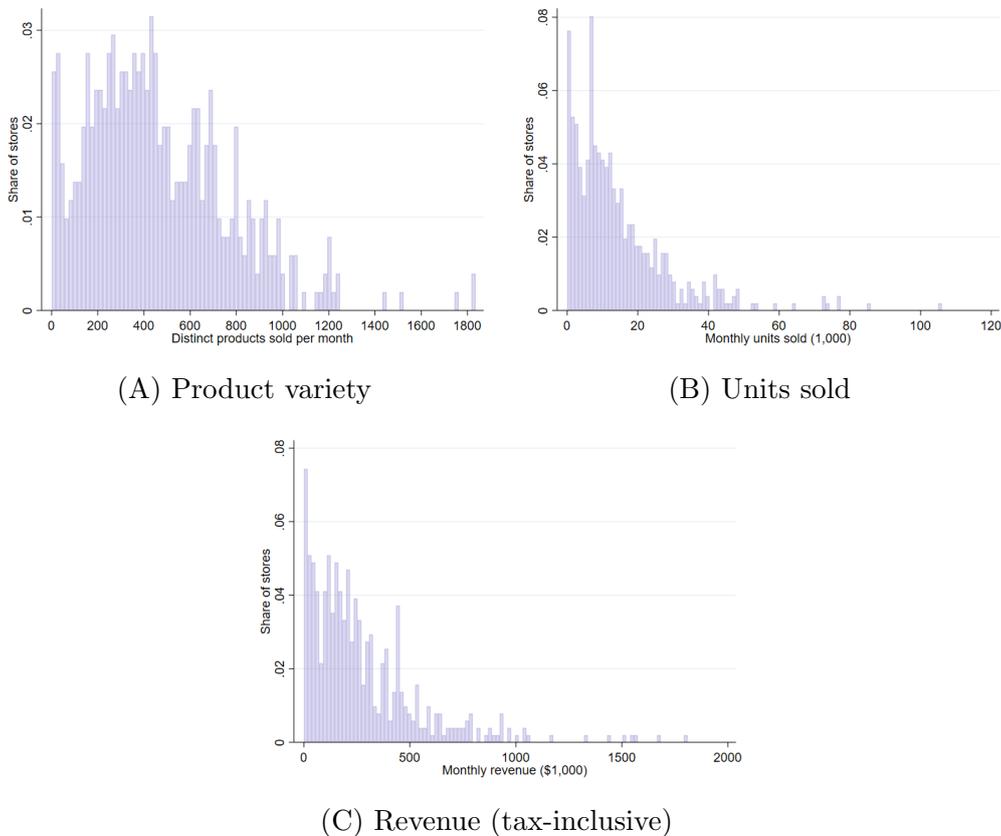


Notes: This figure depicts the flow of cannabis products, from left to right, as they move through the supply chain. Only licensed producers are permitted to cultivate and harvest cannabis plants; producers can only sell to licensed processors, who in turn are permitted to process products; only processors can sell finished products at wholesale to retailers; licensed retailers can sell finished products to end consumers. An establishment can jointly hold producer and processor licenses, so the overwhelming majority of upstream establishments hold both licenses (i.e. producer-processors). Retailers may not hold a producer or a processor license and vice versa. As a result, production and retail activities are legally separated.

A.2 Cannabis retail stores

Store characteristics Figure A.3 shows the distribution of store-level monthly averages for various store characteristics. Panel A shows the average number of distinct products sold per month across stores. A 1.0 gram package and a 2.0 gram package of Sunset Sherbert usable marijuana (i.e. unprocessed dried flower) produced by Northwest Harvesting Co are examples of distinct products in our data. The average store in our sample sells 491 distinct products per month (median: 419). However, Panel A reveals substantial variation across stores in our sample, with values ranging from as low as 13 to a maximum of 1,833 products per month. Panel B reports the average units sold per month across stores (in thousands). The average store in our sample sells 14,391 units per month (median: 10,905). As is the case with product variety, there is large variation in units sold across stores. Stores at the 1st percentile sell 287 units per month, while those at the 99th percentile sell 72,826 units per month. Panel C shows the distribution of tax-inclusive monthly revenue across stores. The average store generates \$276,842 revenue per month (median: \$205,377). Again, revenue varies across stores: stores at the 1st percentile generate \$3,765 while those at the 99th percentile generate \$1,447,000 per month.

Figure A.3: Distribution of monthly averages across stores



Notes: The figures show the distribution of store-level average statistics across all stores in our sample. Panel A presents the distribution of the average monthly number of unique products sold. Panel B displays the distribution of the monthly average units sold. Panel C illustrates the distribution of average monthly sales revenue across stores.

Variable cost structure for cannabis retailers To ascertain the variable cost structure for cannabis retailers, we use aggregate payroll data on cannabis retailers from the Washington State Employment Security Department (ESD) and High Peak Strategy, an economics consulting firm. The ESD collects data on employment and wages in industries covered by unemployment insurance (95% of U.S. jobs). The data spans the years 2018-2020. Table A.1 illustrates that cannabis retailers have a similar variable cost structure as other retail industries studied in the literature. Renkin et al. (2022), for example, find that for U.S. grocery stores, COGS accounts for 83% of variable costs. Note that in most retail settings, cost of goods sold (COGS) and labor cost together account for 99% of variable cost while other expenditures like packaging and transport typically make up less than 1% of variable cost (Renkin et al., 2022).

Table A.1: COGS and the labor share of costs for cannabis retailers

Year	Average expenditure		Variable cost share	
	Labor	COGS	Labor	COGS
2019	\$370,897	\$1,138,239	0.25	0.75
2020	\$407,273	\$1,597,671	0.20	0.80

Notes: This table compares average annual labor expenditure and COGS expenditure for cannabis retail stores in Washington state for the years 2019-2020. Aggregate payroll data on cannabis retailers is from the Washington State Employment Security Department and High Peak Strategy (2019-2020). Labor expenditure equals total wages divided by the number of active stores. Stores with missing UI data are excluded from total wages and establishment counts. COGS is the average annual wholesale expenditure for cannabis retailers in the estimation sample. Wholesale purchases from processor-only licenses are included.

The geography of wholesale costs Table A.2 shows the percentage of retailers’ wholesale costs in relation to a producer’s geographic location. Column 1 shows that only about 5% of retailers’ wholesale expenditures go to producers located in the same city as the retailer. Column 2 shows that less than 15% goes to producers in the same county as the retailer. For Column 3, we sort counties into their respective 3-digit zip codes (retailers are located in 14 3-digit zip codes compared to 37 counties). Column 3 shows that less than 16% of wholesale cost goes to producers located in the same 3-digit zip code. Next, we sort counties into three regions (west, central, east), defined by well-established topographic and economic boundaries. Column 4 shows that about 64% of wholesale sales go to retailers in a different region than the producer. Column 5 looks at the subset of establishments located in the west and east regions of the state, thus dropping producers in the central region. The east and west regions are non-contiguous and are located on opposite sides of the state. For establishments located in these two regions, nearly 23% of wholesale sales go to retailers located in the other region, that is to say, retailers on the opposite side of the state. Because the majority of retail establishments are located in the west and east regions, this share amounts to 21.4% of all of wholesale expenditures in the industry. Taken together, the results from Table A.2 illustrate that there is no home bias in wholesale cannabis purchases.

Table A.2: Share of retailers’ wholesale costs by geographic proximity

	(1)	(2)	(3)	(4)	(5)	(6)
	Same city	Same county	Same 3-digit zip code	Same region	Non-contiguous region	Same state
Percent of wholesale expenditure	7.49%	17.79%	18.06%	65.95%	19.06%	100%

Notes: This table shows the share of retailers’ wholesale expenditure according to wholesalers’ geographic proximity. The shares are based on 11.8 million unique wholesaler-retailer-product-month observations from March 2018 through December 2021. Retailers are located in 14 3-digit zip codes and 35 counties. Region groups counties into three categories: west, central, or east. Data from Top Shelf Data.

A.3 Geography of crimes

Table A.3 shows geographic characteristics of reported crime incidents during our sample period. Column 1 illustrates that approximately half of the crime incidents are in King county, the state’s most populous county, while only 22 percent of all cannabis stores are located in King county (column 4). Approximately two-thirds of crimes are in the Seattle MSA, which contains 38 percent of all stores. The five most populous cities in the state have

55 percent of the crimes but only 27 percent of all stores. Taken together, the relatively high prevalence of crime in these areas is consistent with evidence suggesting that retail crime may be more prevalent in urban areas (Lopez et al., 2023; National Retail Federation, 2024).

Table A.3: The geographic distribution of retail crimes in cannabis

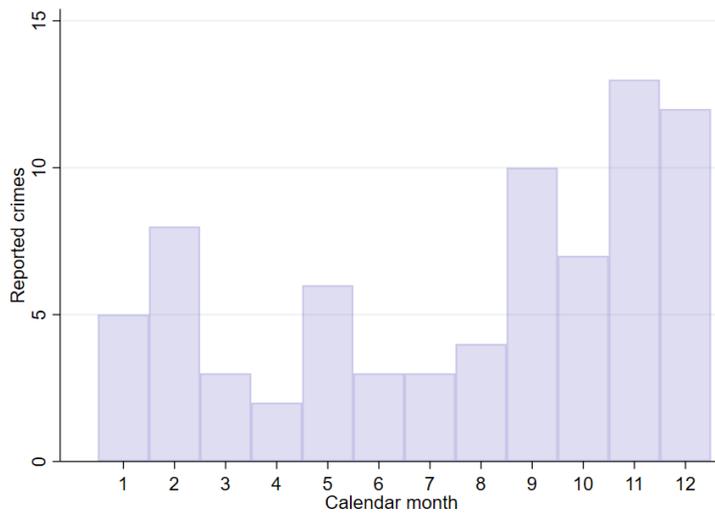
	(1) Crimes	(2) Share of total crimes	(3) Stores	(4) Share of total stores
King county	38	0.51	110	0.22
Not King county	37	0.49	401	0.78
Seattle MSA	50	0.67	193	0.38
Not Seattle MSA	25	0.33	318	0.62
5 most populous cities	41	0.55	136	0.27
Not 5 most populous cities	34	0.45	375	0.73

Notes: The table shows descriptive statistics for the location of reported retail crime incidents at cannabis stores during our sample period. Seattle MSA comprises King, Pierce, and Snohomish counties. The five most populous cities are Seattle, Bellevue, Tacoma, Vancouver, and Spokane.

A.4 Seasonality of crimes

Figure A.4 illustrates the seasonality in crime incidents within our sample. Each bar represents the number of crimes that occurred in the respective calendar month, with 1 indicating January and 12 indicating December. The figure shows that crime occurs throughout all months but that rates tend to be higher towards the end of the year.

Figure A.4: Reported crimes at cannabis retailers by calendar month



Notes: This figure illustrates the seasonality in treatment timing within our sample. Each bar represents the number of crimes that occurred in the respective calendar month, with 1 indicating January and 12 indicating December.

B Establishment-level indexes

This section provides more detail on the store-level indexes used in our analysis.

B.1 Price indexes

Our empirical analysis uses traceability data provided by the data analytic firm Top Shelf Data (TSD), which ingests the raw tracking data from the Washington state Liquor and Cannabis Board (LCB) and matches it with additional product information. Note that the raw tracking data from the LCB includes each product’s Stock Keeping Unit (SKU), but TSD does not report this. Instead, each product is identified by a unique combination of five elements: retailer-producer-category-product name-unit weight. For products with no unit weight (such as liquid edibles), the first four elements identify the product. TSD then calculates the average price of product i at retail store j in month t as

$$P_{i,j,t} = \frac{TR_{i,j,t}}{TQ_{i,j,t}}. \quad (1)$$

where $TR_{i,j,t}$ is the revenue from product i at retailer j in month t , and $TQ_{i,j,t}$ is total quantity.

To construct store-level price indexes, we employ a two-step process similar to that used by Renkin et al. (2022). In the first step, we use $P_{i,j,t}$ to construct a geometric mean of month-over-month changes for product subcategory c at store j :

$$I_{c,j,t} = \prod_i \left(\frac{P_{i,j,t}}{P_{i,j,t-1}} \right)^{\omega_{i,c,j,y(t)}} \quad (2)$$

where each subcategory is a unique category-unit weight combination.¹ For example, usable marijuana is a category, whereas 1.0 gram usable marijuana and 2.0 gram usable marijuana are separate subcategories. Following Renkin et al. (2022), the weight $\omega_{i,c,j,y(t)}$ is the share of product i in total revenue of subcategory c in establishment j during the calendar year of month t .²

In the second step, we aggregate across subcategories to get the price index for store j in month t :

$$I_{j,t} = \prod_c I_{c,j,t}^{\omega_{c,j,y(t)}}. \quad (3)$$

Similar to the first step, the weight $\omega_{c,j,y(t)}$ is the share of subcategory c in total revenue in

¹Since unit weight is a major component of cannabis product differentiation (akin to volume in beverage sales), the majority of sales contain information on unit weight. Therefore, in the first step of the establishment index, we choose to aggregate at category-unit weight level rather than the category level.

²As pointed out by Renkin et al. (2022), price indexes are sometimes constructed using lagged quantity weights. Since product turnover can be high in retail settings—and cannabis is no exception—lagged weights would limit the number of products used in constructing the price indexes. Thus, contemporaneous weights are used.

store j during the calendar year of month t . The store-level inflation rate is then simply the natural logarithm of the index

$$\pi_{j,t} = \ln I_{j,t} \quad (4)$$

B.2 Quantity indexes

The quantity indexes are constructed the same way as the price indexes. The only difference is in the first step

$$I_{c,j,t} = \prod_i \left(\frac{Q_{i,j,t}}{Q_{i,j,t-1}} \right)^{\omega_{i,c,j,y(t)}} \quad (5)$$

where $Q_{i,j,t}$ is the quantity sold of product i at store j in month t . The index weights are otherwise identical to those from the price index.

B.3 Wholesale cost indexes

The retail scanner data provided by Top Shelf Data contains granular information (at the store-product-month level) on the wholesale prices that retailers pay for each product. We use this information to construct wholesale cost indexes. The wholesale cost indexes are constructed similar to the price indexes. TSD calculates the average monthly wholesale price at the SKU-level as

$$W_{i,j,t} = \frac{TE_{i,j,t}}{TQ_{i,j,t}}. \quad (6)$$

where $TE_{i,j,t}$ is the total expenditure on product i at retailer j in month t , and $TQ_{i,j,t}$ is total quantity purchased at wholesale. In the first step of the establishment-level cost index, we construct a geometric mean of the month-over-month changes in wholesale price for product subcategory c at store j :

$$I_{c,j,t} = \prod_i \left(\frac{W_{i,j,t}}{W_{i,j,t-1}} \right)^{\omega_{i,c,j,y(t)}} \quad (7)$$

where the subcategories are the same as before. In contrast to the price indexes, $\omega_{i,c,j,y(t)}$ is an expenditure weight equal to the share of product i in the wholesale expenditure of subcategory c at store j in the calendar year of month t . Unlike with wholesale prices, information on retailers' wholesale quantity purchased is reported by producer-processors in a separate dataset provided by Top Shelf Data. Since producer-processors sometimes use different product naming conventions than retailers, using the producer-processor data on quantities results in a sizable reduction in the number of products for which retailers' expenditures can be matched. To circumvent this issue, we use the monthly retail quantity sold as a proxy for the monthly quantity purchased by retailers at wholesale. We then calculate the wholesale expenditure for each store-product-month as the monthly wholesale price multiplied by the monthly retail quantity sold. Finally, we sum across months to obtain annual store-product expenditures which serve as inputs for the store-product expenditure weights (the same logic applies for the subcategory weights). By using retail quantity sold to construct the expenditure weights,

we assume that at the product level retailers purchase wholesale quantities in proportion to their retail quantity sold. We view this as unproblematic, especially over longer time horizons as is the case with the annual expenditure shares.

Note that by calculating the expenditure weight based on annual (rather than monthly) expenditure shares, we may not capture short-term wholesale substitution patterns on the part of retailers. To check this, we also construct the index using monthly expenditure shares. This increases the variance of the indexes but does not change our main estimation results. Therefore, we use the annual expenditure weights for our preferred specification.³

In the second step, we aggregate across subcategories to get the wholesale cost index for store j in month t :

$$I_{j,t} = \prod_c I_{c,j,t}^{\omega_{c,j,y(t)}}. \quad (8)$$

Similar to the first step, the weight $\omega_{c,j,y(t)}$ is the share of subcategory c in total wholesale expenditure in store j during the calendar year of month t . The store-level wholesale cost index is then simply the natural logarithm of the index, $\ln I_{j,t}$.

B.4 Descriptive statistics of store level indexes

Table A.4 provides descriptive statistics for our indices. The price and wholesale cost indices are centered around zero and have means close to zero. Standard deviations for these indices range from 0.02 to 0.03. The quantity index has a larger standard deviation than the other indices. This is a common characteristic of quantity indices constructed from store-level scanner data, and is similar to what is found elsewhere in the literature (see e.g. Renkin et al., 2022). The distributions of the indices are presented in Figure A.5.

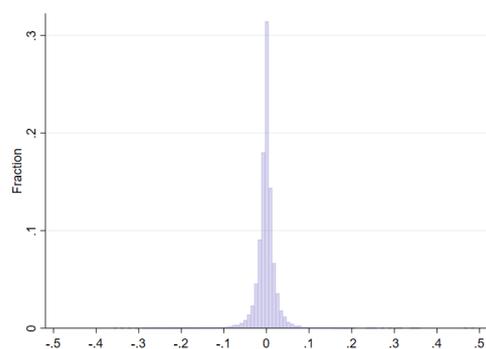
Table A.4: Dependent variable descriptive statistics

	Mean	SD	Median
Log price index	-0.001	0.030	0
Log quantity index	-0.008	0.216	-0.013
Log wholesale cost index	0.00003	0.021	0

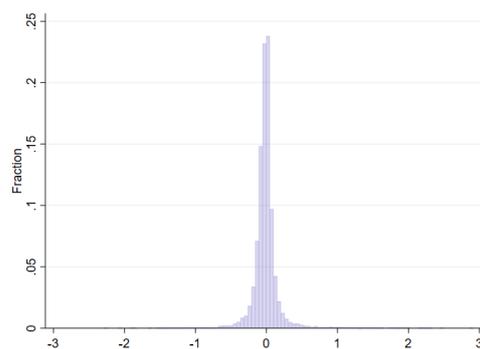
Notes: This table presents the mean, standard deviation, and median of our main dependent variables: Log price index, Log quantity index, and Log wholesale cost index.

³Results using monthly weights are available on request.

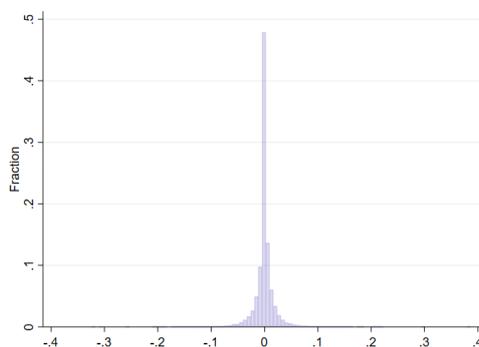
Figure A.5: Distributions for store-level indexes



(A) Price



(B) Quantity sold



(C) Wholesale cost

Notes: The figures show the distribution of the main dependent variables over the sample period. Panel A depicts the monthly store-level price index. Panel B illustrates the monthly store-level quantity index. Panel C displays the monthly store-level wholesale cost index. For visual clarity, extreme outliers are omitted from the figures, although they are included in our regressions.

C Additional heterogeneity analysis and descriptive statistics

C.1 Independent versus chain stores

Table A.5: Cannabis chain store characteristics

	(1) Independent	(2) Chains
Unit price (in dollars)	26.74 (4.29)	26.78 (4.58)
Units sold per month	13,421 (13,068)	16,161 (15,399)
Monthly revenue (in dollars)	257,473 (252,444)	312,154 (301,285)
Unique products per month	460 (379)	553 (391)
N	330	181

Notes: The table reports descriptive statistics for chains and independent stores using the sample of active cannabis retail stores as defined in the main paper. Standard deviations are in parentheses. The variables reported are: unit price, average store quantity sold per month, average store revenue per month, and average number of distinct products sold per month.

In Section 5.2, we distinguished between small and large chains, with the former defined as stores belonging to firms with 1-2 stores and the latter corresponding to three or more stores. Table A.5 reports descriptive statistics for the different store types, based on all 511 active cannabis retailers.

In Table A.7 we report alternative specifications for our heterogeneity analyses. In this specification we define independent stores as stores belonging to a single-store firm, and chain stores as stores at firms with two or more stores. Column 1 and 5 show that effects at independent stores tend to be smaller than the effects at small chains found in Section 5.2. Moreover, Columns 2 and 6 show a larger the effect at chains than in those found in Section 5.2. This redistribution of treatment effects reflects that stores in small chains (i.e. two-store chains) have a large treatment effect that is similar in magnitude to independent stores. This points to small chain stores facing similar competitive pressures as independent stores, and validates our definition of independent and chain stores used in our main heterogeneity analysis in Section 5.2.

C.2 Market concentration

Table A.6 shows the descriptive statistics for our store-level market concentration measure, as defined in our main text.

Table A.6: Market concentration descriptive statistics

	(1) Mean	(2) Median	(3) Std
Victimized HHI	0.14	0.09	0.13
All HHI	0.31	0.22	0.24

Notes: The table shows descriptive statistics for our market concentration measure, calculated on the store-level. The first row only considers victimized stores. the second row considers all stores within the Washington State cannabis industry.

In Section 5.2, we found large effects at stores in markets with low concentration, and no effect at markets with high concentration. One possible reason for this finding is that low concentration markets are generally located in urban areas, where labor costs (and hence the security cost shock) may be higher compared to rural areas. To test this, we use an alternative definition of low and high concentration markets that accounts for urban-rural heterogeneity. We proceed as follows. First, we categorize as urban stores that are located in the four largest and most densely populated cities (Seattle, Tacoma, Bellevue, and Spokane), and consider all other stores as rural. Next, we calculate the median HHI for stores in the urban subsample and the median HHI for stores in the rural subsample. We split the urban subsample into high concentration and low concentration parts, and do the same for the rural subsample. This leaves us with four subsamples: urban low concentration, urban high concentration, rural low concentration, and rural high concentration. Finally, we pool the high concentration subsamples into a single subsample, and the low concentration subsamples into another subsample. Our low concentration subsample thus contains stores in markets considered low concentration in an urban setting, but also stores in markets considered low concentration in a rural setting. We then estimate our main regression equation separately for the high concentration and the low concentration subsamples. Columns 3-4 show that at victimized stores, six months after a crime the difference in effect sizes between low and high concentration markets is 0.01 while the difference in effect sizes in Section 5.2 is 0.021. Thus, accounting for the urban-rural divide appears to explain the comparatively large effect size for low concentration markets found in Section 5.2. Similarly, Columns 7-8 show an equalization of the effect sizes for rivals in low- and high-concentration markets. Without accounting for the urban-rural divide, the difference in effect sizes at $t+6$ is 0.015 (Table 3, while accounting for the rural-urban divide results in a slightly higher effect in high concentration markets.

To summarize, for both victimized and rival stores alike, urban-rural heterogeneity appears to explain much of the heterogeneous effects for low and high concentration markets,

particularly at higher lags.

Table A.7: Heterogeneity analysis robustness checks

Cumulative price effect	Victimized				Rivals			
	(1) Indep. stores	(2) Chain stores	(3) Low concentration	(4) High concentration	(5) Indep. stores	(6) Chain stores	(7) Low concentration	(8) High concentration
t	0.017 (0.012)	0.0014 (0.0033)	0.013 (0.010)	0.0018 (0.0028)	0.0011 (0.0045)	0.0022 (0.0023)	0.00063 (0.0025)	0.0034 (0.0036)
$t + 2$	0.030** (0.015)	0.0058 (0.0055)	0.025** (0.012)	0.0061 (0.0056)	0.0026 (0.0050)	0.0077* (0.0043)	0.0028 (0.0039)	0.011* (0.0060)
$t + 4$	0.035** (0.016)	0.0061 (0.0083)	0.030** (0.014)	0.0052 (0.0084)	0.0099* (0.0055)	0.019*** (0.0069)	0.011 (0.0067)	0.024*** (0.0081)
$t + 6$	0.023 (0.020)	0.018*** (0.0067)	0.025 (0.017)	0.015** (0.0068)	0.010 (0.0071)	0.013* (0.0080)	0.012 (0.0083)	0.015* (0.0082)
\sum Pre-event	0.011 (0.0090)	-0.0075 (0.0072)	-0.0055 (0.0096)	0.0047 (0.0063)	-0.00050 (0.011)	0.0065 (0.0069)	-0.0019 (0.0073)	0.010 (0.0082)
N	11,496	18,218	14,501	15,213	28,626	29,535	15,673	14,954

Notes: Each column shows the cumulative treatment effects on store price levels for different subsamples zero, two, four and six months after a crime, along with the sums of pre-treatment coefficients. Coefficients are interpretable as percentage increases in outcome levels relative to the month before a crime incident. The first four columns use victimized stores as the treatment group, and the last four columns consider rival stores. Columns 1 and 5 show effects for independent stores (owned by firms running one store only), while columns 2 and 6 show effects for stores owned by firms running at least two stores. For the other columns, the Herfindahl-Hirschman Index, conditional on a store being located in a rural or urban area, is calculated for the market around each store (including non-treated stores) within a 5-mile radius. The sample is then split according to the median market concentration in rural and urban areas, respectively. Columns 3 and 7 show effects for treated stores in rural and urban markets with below median concentration, and columns 4 and 8 for treated stores in markets with above median concentration. Standard errors of the sums are clustered at the store level and shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

D Inclusion criteria for our control group and spatial spillovers

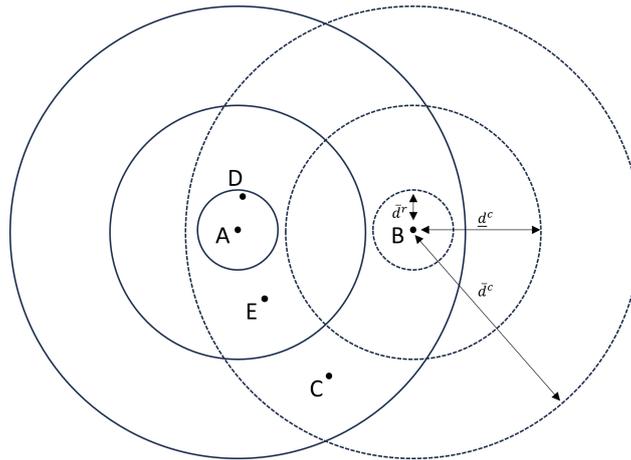
This section elaborates on the inclusion criteria for our control group, with a special emphasis on the geographical criteria.

D.1 Stacked DiD with spatial treatment

To provide a clearer intuition behind our inclusion criteria for defining clean control stores, we provide an illustrative example in Figure A.6. The figure depicts five stores: A and B are both victimized while stores C, D, and E are not victimized. Three concentric rings surround each victimized store: the innermost ring denotes the boundary for rival stores and has a radius of \bar{d} . The other two rings, with radius \underline{d}^c and \bar{d}^c , denote the boundaries for control candidates.

Store C is a control candidate for both A and B while store E is a control candidate for B only. Store D is at the same time a control candidate for B and a rival of A. If the event windows for stores A and B overlap, D is disqualified from the control group for B. However, store D qualifies as a clean control for store B if the event windows for stores A and B do not overlap. This reflects our assumption that store D's contamination from store A's crime incident is zero outside the treatment window. Similarly, an earlier-treated store (e.g. A) would qualify as a control store for a later-treated store (e.g. B) if the earlier-treated store is a control candidate and the event windows do not overlap.

Figure A.6: Clean controls



Notes: The figure depicts the locations of four stores (A, B, C, D) in relation to each other. Stores A and B are victimized while stores C and D are not. Concentric rings of diameter \underline{d} and \bar{d} denote the boundaries defining clean control stores in our setting.

Using the intuition from Figure A.6, we can formally express the inclusion criteria for clean controls in our stacked DiD design. Let H denote the set of all stores in our sample

and J the set of all victimized stores, chronologically labeled $j = 1, \dots, n$. Define $R_j = \{h \neq j \in H : d_{h,j} < \bar{d}^r\}$ as the set of rival stores related to victimized store j , with $R = \bigcup_{j=1}^n R_j$ representing the set of all unique rival stores. t_j denotes the month a store j is treated (i.e., the month of the crime incident in sub-experiment j). w is the size of the event window (in months), $d_{h,j}$ is the geodesic distance between stores j and h (in miles).

Store i qualifies as a clean control and is included in the sub-experiment corresponding to the crime incident at store j if the following conditions jointly hold:

IC.1: $i \in O_j$ where $O_j = \{h \in H : \underline{d}^c \leq d_{h,j} \leq \bar{d}^c\}$

IC.2: $i \notin T_{j-}$, where $T_{j-} = \{h \in J : (t_h + w > t_j) \vee (t_h - w < t_j)\}$

IC.3: $i \notin D_{j-}$, where $D_{j-} = \{h \in R : (t_h + w > t_j) \vee (t_h - w < t_j)\}$

Condition IC.1 requires that store i be located between \underline{d}^c and \bar{d}^c miles of store j . This geographic criterion designates candidate control stores for each treated store, e.g., stores A, C, D, and E for store B in Figure A.6. Condition IC.2 ensures that store i is either never victimized or, if it is victimized, its treatment event window does not overlap with that of victimized store j . This is the standard stacked DiD inclusion criterion that prevents unclean comparisons between early- and late-treated stores. In our example, this criterion excludes store A from the set of clean control stores for B and vice versa if the event window of A and B overlaps. Condition IC.3 requires that store i must not be a rival store of any victimized store whose event window overlaps with that of store j .⁴ This condition would disqualify store D as a clean control for store B if the treatment windows of stores A and B overlap. However, IC.2 and IC.3 ensure that stores D and A qualify as a clean control candidate for store B if store B's event window does not overlap with that of store A.

Note that store i may satisfy IC.1-IC.3 for some crime incidents but not for others. As long as store i satisfies all three criteria for a given crime incident, then it will be included in that sub-experiment. Accordingly, store i may qualify as a clean control for multiple sub-experiments that overlap in calendar-time. This is a common feature of the stacked DiD estimator and it implies that certain store-month observations may be recycled and appear in multiple sub-experiments (Wing et al., 2024). For example, in Figure A.6, if stores A and B are both victimized in the same period, then store C serves as a clean control for both sub-experiments (assuming the other conditions are satisfied). Store C is then included in each of the sub-experiment-specific datasets, meaning the stacked dataset will contain duplicate observations for store C for the calendar months corresponding to that particular event window.

⁴Note that rival treatment is only assigned according to a store's first nearby crime incident, implying that each rival store is assigned to only one treatment cohort and that multiple treated rival stores can serve as clean control stores for subsequent crimes even if the event window overlaps with store j . This cannot happen to victimized stores, as we exclude multiple victimized stores before the second treatment.

D.2 Spatial spillovers

In defining our control group, we face a tradeoff. On the one hand, we want to ensure that stores in the control group are in markets that are comparable to (and hence near) the treated group. At the same time, if there are spillovers from the treated group to nearby stores, then including nearby stores in the control group leads to biased estimates (Muehlegger and Sweeney, 2022).

These spillovers could arise from various factors. One concern is strategic complementarity in pricing. Muehlegger and Sweeney (2022) show that in the presence of imperfect competition, the strategic response of (untreated) competitors may disqualify them as a valid control group. In practice, this implies that the price a firm sets is a function of not just its own costs, but also those of its rivals. We provide a detailed discussion and analysis of this issue in section E.

Spillovers can also arise if consumers substitute demand out of victimized stores and into nearby competitors. While demand substitution is distinct from strategic complementarity in pricing, at a mechanical level, both entail consumers choosing between stores from the same choice set. Therefore, we consider the 30-mile boundary sufficiently distant to prevent demand substitution from biasing our estimates.

Finally, spillovers can occur if crime incidents induce a cost shock at rival stores. These spillovers could emerge through increased precautionary security measures or higher expected future insurance expenditures, similar to the responses at treated stores. While the geographic scope of such spillovers is difficult to quantify, it is plausible that stores located close to a victimized store are more likely to face increased costs, e.g. due to local market characteristics. Therefore, we expect these spillover effects to operate within a similar geographic radius as strategic complementarity in pricing and demand substitution. Additionally, if control stores experience rising marginal costs, we likely underestimate the treatment effects.

E Strategic complementarity in pricing

In this appendix section, we discuss the implications and potential issues arising from strategic complementarity in pricing in our setting. Furthermore, we present estimation results measuring the extent of strategic pricing in the Washington state cannabis industry.

E.1 Theoretical framework

We follow the framework of Muehlegger and Sweeney (2022) and consider the pass-through of a tax (or input cost shock) τ onto the price of firm j . Firm j sets the profit-maximizing price p_j and faces tax-inclusive marginal costs α_j . Each firm in the market can have a different exposure to the tax, with $\frac{\partial \alpha_j}{\partial \tau}$ capturing the marginal unit tax rate faced by firm j . In oligopolistic markets, the price a firm sets is a function of not just its own costs, but also

those of its rivals. The pass-through of the tax onto firm j 's price can thus be decomposed as a direct (own-cost) and an indirect (competitors' cost) effect:

$$\frac{\partial p_j}{\partial \tau} = \frac{\partial p_j}{\partial \alpha_j} \frac{\partial \alpha_j}{\partial \tau} + \sum_{i \neq j} \frac{\partial p_j}{\partial p_i} \frac{\partial p_i}{\partial \alpha_i} \frac{\partial \alpha_i}{\partial \tau} \quad (9)$$

where $\frac{\partial p_j}{\partial p_i}$ is firm j 's best response to a change in firm i 's price.⁵ Consequently, in the presence of imperfect competition, the strategic response of (untreated) competitors may disqualify them as a valid control group. It is therefore important to quantify the size and geographic scope of strategic complementarity in prices.

E.2 Quantifying strategic complementarity in prices

To identify the scope of strategic complementarity in prices, we follow the industrial organization literature that measures the pass-through of cost shocks and taxes. In particular, we build on the approach of Hollenbeck and Uetake (2021), who use similar data to evaluate the optimal cannabis sales tax. A major advantage of this approach is that, because we observe wholesale unit prices, we can directly measure how changes in unit cost are passed through to prices.⁶ In addition to stores' own wholesale unit costs, we also observe the wholesale unit costs of their competitors. By relating stores' prices to competitors' cost changes, we can measure the effect of competitors' cost-induced price changes, i.e. strategic complementarity in prices. Moreover, we can test whether this effect is a function of the geographic distance between stores. We use the results of this analysis to define unaffected local markets (a clean control inclusion criterion in Section 4).

To investigate the geographic scope of strategic complementarity of prices, we sort competitors into 5-mile bins and calculate average wholesale unit price for each store-product-month-bin. We specify a model at the store-product-month level that relates a store-product's retail price to (i) the wholesale unit price and (ii) the average wholesale unit price paid by stores in each distance bin. By including both own costs and competitors' costs, we capture the total effect (i.e. own-cost and strategic price response) of an aggregate unit cost shock on stores' prices. Since cannabis transaction data is publicly available, stores have full information on competitors' unit costs and prices updated on an almost weekly basis. Therefore, we focus on contemporaneous changes in costs and prices. This is in line with the pass-through literature from other industries (see e.g. Hollenbeck and Uetake, 2021; Muehlegger and Sweeney, 2022; Conlon and Rao, 2020; Miller et al., 2017). We estimate the following

⁵For ease of exposition we consider competition in prices. Muehlegger and Sweeney (2022) show that this framework extends to a broad class of oligopolistic settings.

⁶Wholesale costs are typically estimated from supply-side first order conditions. For similar approaches, see, for instance, Muehlegger and Sweeney (2022); Ganapati et al. (2020) who use variation in energy input costs to estimate the price pass-through of a hypothetical carbon tax or Miller et al. (2017) who estimate the pass-through of carbon pricing in the portland cement industry.

model in first-differences:

$$\Delta p_{i,j,t} = \rho \Delta w_{i,j,t} + \sum_{r=1}^R \beta_r \Delta w_{i,r(j),t} + \Delta \gamma_t + \Delta \varepsilon_{i,j,t}, \quad (10)$$

where $p_{i,j,t}$ is the average price (in dollars) of product i sold at store j in month t , $w_{i,j,t}$ is the average wholesale price that retailer j pays for product i in month t , $w_{i,r(j),t}$ is the average wholesale price that competitors pay for product i in month t , and γ_t is the year-month FE. In our baseline specification, we set $R = 9$ ($R > 9$ does not meaningfully affect estimates but changes the sample size and standard errors). We cluster standard errors at the store level to allow for autocorrelation in unobservables within stores.

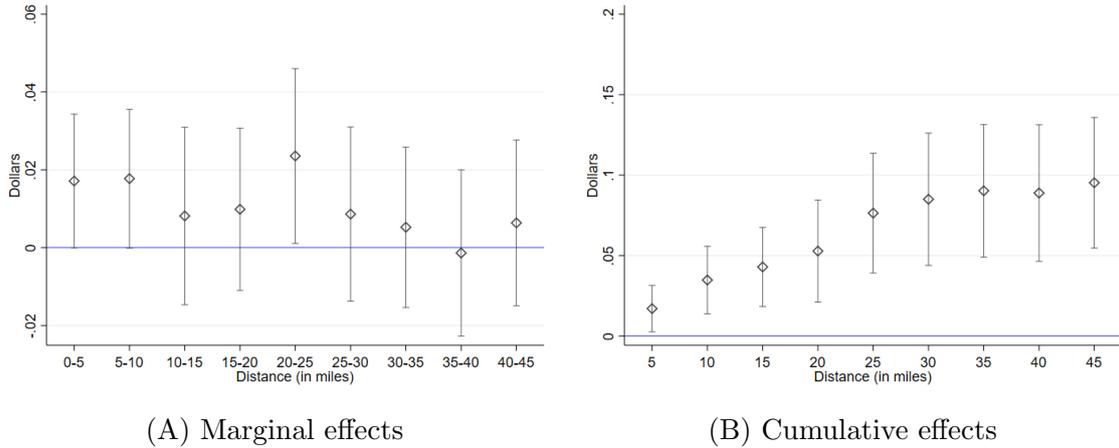
The effect of an aggregate change in unit costs on store j 's prices comprises two parts. The first is the own cost pass-through rate, ρ , i.e. the increase in retail unit price at store j from the increase in store j 's own wholesale unit cost. The second part is β_r which measures the pass-through of wholesale unit costs at competing stores in bin r to unit prices at store j . This is equivalent to the strategic price response between store j and competing stores in bin r .

In Figure A.7 Panel A, we report estimated pass-through rates of competitors' unit costs, β_r , from our baseline specification. The estimates differ across bins with the largest effects in the 5-10 mile and 20-25 mile bins. The fact that the effect fluctuates with distance could reflect commuting patterns, with the average daily distance travelled in Washington state ranging from less than 20 miles in some counties to more than 70 in others (Axios, 2024). Nevertheless, at the 25-30 mile bin, effects shrink and remain close to zero for three consecutive bins.

The β_r estimates in Panel A can be interpreted as marginal effects in that they measure the additional effect on store j 's prices of increasing the geographic scope of an aggregate change in costs by another 5 miles. While this is informative about the geographic scope of strategic complementarity in cannabis prices, quantifying the actual effect on prices of an aggregate change in costs requires summing the marginal effects $\sum_{r=1}^R \beta_r$ up to a given distance bin R . The sum can be interpreted as the effect on store j 's prices of an aggregate change in costs that affects all stores up to a given distance (while holding store j 's costs constant). We report these sums at increasing distances in Panel B of Figure A.7. Panel B further highlights that sensitivity to competitors' costs increases up to the 30-mile mark before plateauing thereafter. This aligns with a growing literature showing that the scope of cost shocks matters and that aggregate (i.e market-wide) cost shocks elicit a larger strategic price response than idiosyncratic or highly localized shocks (Muehlegger and Sweeney, 2022).

In addition to our main specification, we report several variants of our pass-through regression in Table A.8. In Column 2, we add controls (county-level average wages, the county-level unemployment rate, and the home price index at the three-digit zip code level) to absorb variation in retail cannabis prices due to local business cycles or changes in house

Figure A.7: The pass-through of competitors' wholesale unit costs to own unit prices



Notes: Panel A shows estimated coefficients β_r for $r \in [1, 9]$ obtained from the pass-through regression (equation 10). Coefficients in Panel A are interpretable as the effect (in dollars) on store j 's retail unit price from a \$1 increase in wholesale unit costs affecting all stores in distance bin r . Panel B shows cumulative sums of coefficients $\sum_{r=1}^R \beta_r$ for $R \in [1, 9]$. Coefficients in Panel B are interpretable as the effect on store j 's retail unit price of a \$1 increase in wholesale unit cost affecting all stores up to r miles away. The figure shows 90% confidence intervals of the sums based on SE clustered at the store level. Data: Top Shelf Data, March 2018 through December 2021.

prices (see Stroebel and Vavra, 2019). In Column 3, we include region-time FE to account for other spatially correlated shocks that may covary with wholesale and retail cannabis prices. We also estimate equation 10 in levels rather than first-differences, with store-product FE (Column 4). In Column 5, we specify equation 10 using the first-difference of logs. This minimizes the influence of outliers and delivers pass-through elasticities instead of pass-through rates.

Columns 1-5 in Table A.8 show that stores' sensitivity to competitors' costs plateaus at the 30-mile mark across all specifications. Moreover, the results indicate that an aggregate cost shock with sufficient geographic scope has non-negligible strategic price effects. When estimated in first-differences (Column 1), a \$1 increase in wholesale unit costs at all stores within a 30-mile radius corresponds to a \$0.09 increase in retail prices solely due to strategic complementarities. When estimated in levels (Column 4), prices increase \$0.47 from a \$1 increase in wholesale cost.

Overall, the results from Figure A.7 and Table A.8 provide suggestive evidence of strategic complementarity in prices for cannabis stores within 30 miles of each other. Increasing the geographic scope of an aggregate cost shock appears to have little additional effect on store prices beyond the 30-mile mark. This suggests that stores located more than 30 miles from a victimized store will not have a strategic price response to the crime-induced cost shock at victimized and rival stores. We therefore view Figure A.7 and Table A.8 as providing supportive evidence for our definition of unaffected local markets from Section 4.

Table A.8: Cumulative pass-through of competitors' unit costs

	Dollars (FD)					
	(1)	(2)	(3)	(4)	(5)	(6)
	Baseline	Controls	Reg. \times time FE	Dollars (levels)	Logs (FD)	Store- level index
Own wholesale cost	1.75*** (0.046)	1.75*** (0.046)	1.75*** (0.046)	1.26*** (0.41)	0.73*** (0.0092)	0.27*** (0.020)
Competitors' wholesale cost						
< 5 miles	0.017* (0.0087)	0.017* (0.0087)	0.017* (0.0087)	0.073** (0.032)	0.0026 (0.0023)	0.0030 (0.0042)
< 10 miles	0.035*** (0.013)	0.035*** (0.013)	0.035*** (0.013)	0.23*** (0.063)	0.0049 (0.0032)	0.0072 (0.0044)
< 15 miles	0.043*** (0.015)	0.043*** (0.015)	0.044*** (0.015)	0.25*** (0.080)	0.0038 (0.0038)	0.0067 (0.0046)
< 20 miles	0.053*** (0.019)	0.053*** (0.019)	0.053*** (0.019)	0.31*** (0.092)	0.0070 (0.0047)	0.0068 (0.0058)
< 25 miles	0.076*** (0.023)	0.077*** (0.023)	0.076*** (0.023)	0.34*** (0.10)	0.010* (0.0056)	0.0072 (0.0069)
< 30 miles	0.085*** (0.025)	0.085*** (0.025)	0.084*** (0.025)	0.47*** (0.13)	0.016*** (0.0056)	0.0075 (0.0076)
< 35 miles	0.090*** (0.025)	0.091*** (0.025)	0.089*** (0.025)	0.51*** (0.14)	0.017*** (0.0055)	0.0015 (0.011)
< 40 miles	0.089*** (0.026)	0.089*** (0.026)	0.088*** (0.026)	0.50*** (0.15)	0.015*** (0.0049)	0.0061 (0.012)
< 45 miles	0.095*** (0.025)	0.096*** (0.025)	0.094*** (0.025)	0.46*** (0.15)	0.021*** (0.0070)	0.0072 (0.013)
<i>N</i>	2,406,037	2,406,037	2,406,037	3,881,127	2,406,037	9,826

Notes: The table reports the estimates of wholesale cost pass-through rates from equation 15. We report estimates for own wholesale cost changes and for average changes in wholesale costs at competitor stores located within 5 miles of the respective store. All specifications control for month-year fixed effects. Coefficients for Columns 1-4 are interpretable as pass-through rates in dollars. Coefficients in Columns 5-6 are interpretable as pass-through elasticities. Standard errors are clustered at the store level and shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Wholesale unit cost pass-through with store-level price and cost indexes

It is worth noting that equation 10 is specified at the store-product-month level of aggregation. There are two reasons for using this disaggregated level in our baseline specification. First, our target parameter is the pass-through of *unit* (i.e. per product) cost to *unit* price. The disaggregation allows us to estimate this parameter by directly relating these two values. In contrast, a regression based on store-level indexes first aggregates products within their respective subcategory and then aggregates across subcategories within a store. While this is preferable when estimating treatment effects at the store level, the two-step aggregation necessarily breaks the direct vertical link between unit cost and unit price. Second, and more importantly, this link is further cleaved by the fact that rivals' cost indexes contain cost changes for products and subcategories that store j may not actually sell. In the extreme case, one might relate costs and prices for adjacent stores that sell non-overlapping baskets of goods and hence do not compete in prices at the product level. As a result, the estimated coefficient for rivals' cost changes—and by extension the own-cost pass-through rate—may be biased when using store-level indexes. Nevertheless, we estimate the pass-through regression using store-level price and cost indexes in Column 6 of Table A.8. The own-cost pass-through elasticity is smaller than the elasticity in Column 5, but highly statistically significant. However, the coefficients on competitors' costs are close to zero and not statistically significant, in line with the idea that competitors' cost indexes may not capture strategic complementarity in prices for the reasons outlined above.

F Alternative definitions of unaffected local markets and rival stores

In this section, we conduct several robustness checks to assess the sensitivity of our estimated effects to these geographical criteria and our definitions of rival stores.

F.1 Heterogeneous spatial treatment and inclusion criteria

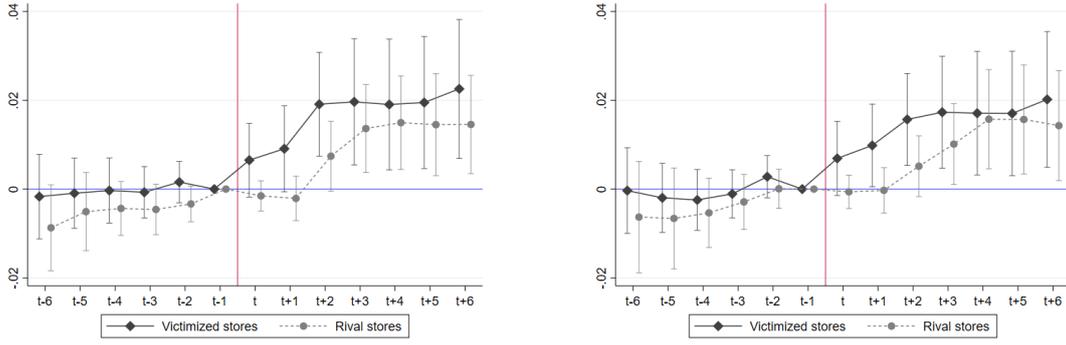
The LCB distributes retail cannabis licenses to counties according to population density, meaning urban areas have a higher concentration of stores than rural areas. Our baseline specification defines the control group in a way that does not account for such heterogeneity. In this subsection, we conduct three robustness checks that allow for heterogeneity in the geospatial density of stores across local markets.

First, we sort stores by their proximity to each victimized store and define clean control candidates as the 150th to 250th closest stores (plus the other clean control criteria discussed in Section D.1). This definition roughly coincides with the average number of candidate controls in our baseline specification, but allows the distances of the inner and outer boundary to vary according to the local market density. We define rivals as the 20 closest stores, which after applying the stacked DiD treatment timing criteria aligns with the average number of rivals in our baseline specification. Panel A in Figure A.8 shows that the rank-based criteria produce similar treatment effect sizes as in our baseline specification.

Second, we revert to distance-based criteria but allow for different inner/outer ring boundaries for stores in urban and rural areas. This allows for heterogeneity across local markets in a restrictive manner. We define urban stores as those located in the largest municipalities in the state: Seattle, Spokane, Tacoma, Vancouver, Bellevue, Kent, Everett, Renton, Spokane Valley (Washington State Office of Financial Management, 2024). For victimized stores in these cities, we set the inner/outer boundary for clean control candidates at 10-30 miles. For all other victimized stores (which we categorize as "rural" for ease of exposition), the inner/outer boundary is 30-60 miles as in the baseline specification. For both urban and rural stores, the definition of rival stores remains 5 miles (results are similar if we reduce the radius for urban rivals). Panel B of Figure A.8 shows results for this specification. Estimated treatment effects are very similar to our baseline specification: four months after a crime, prices at victimized stores are 1.7% higher than the month before the crime while the price effect at rival stores is 1.6%. Moreover, the timing and path of treatment effects are identical to the baseline specification.

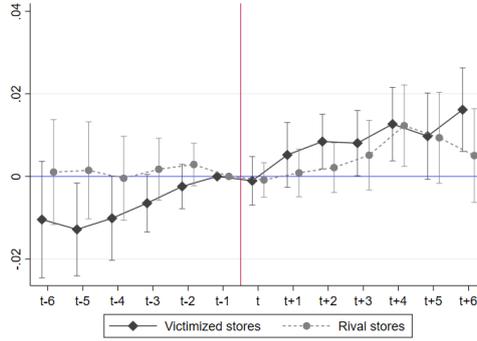
Third, we estimate an alternative specification excluding all stores located in King county, including both treated and potential control stores. King county, which includes the greater Seattle area, is the most populous county in Washington state and has a high concentration of retail crime incidents (see Table A.3). This robustness check addresses concerns that our main results may primarily reflect a King county vs. non-King county comparison. The

Figure A.8: Price effects with spatial heterogeneity



(A) Rank-based criteria

(B) Distance-based criteria



(C) Excluding King County

Notes: Panels A, B and C show cumulative treatment effects for victimized stores and for rival stores, respectively, along with corresponding 90% confidence intervals based on standard errors clustered at the store level. In Panel A, control candidate stores are the 150th to 250th closest stores. Rivals are the 10 closest stores. In Panel B, control candidate stores are 10-30 miles if the victimized store is in an urban area, and 30-60 miles if the victimized store is in a rural area. Rivals are based on a 5-mile radius. In Panel C, we exclude all stores that are located in King County.

estimated effects are slightly attenuated but remain comparable to our main specification and statistically significant four months after a crime for both treatment groups. Taken together, the results from Figure A.8 show that our main results hold when we account for geospatial heterogeneities across urban and rural areas.

F.2 Treatment spillovers from other crime incidents

As discussed above, a store must be at least 30 miles from a victimized store to qualify as a clean control in a sub-experiment. However, unlike victimized and rival stores, this distance restriction does not apply to stores between \bar{d}^r and \underline{d}^c from other victimized stores with overlapping event windows. Thus, stores within 30 miles of a soon-to-be-victimized or recently victimized store can still serve as clean controls if they meet all other criteria. For instance, in Figure A.6, store E is a clean control store for store B, even if the event windows of store A and B overlap. In this section, we investigate the sensitivity of our results to applying more restrictive clean control criteria to address potential spillover effects on these

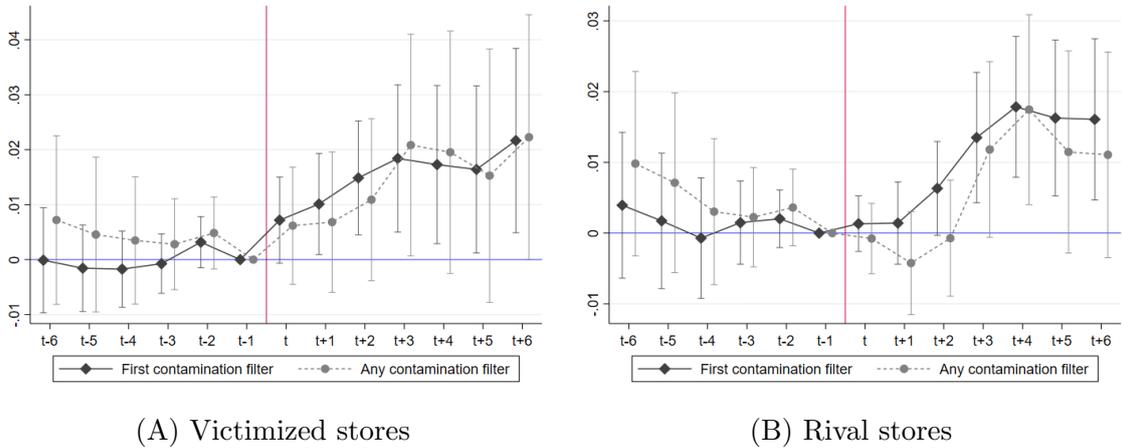
stores.

Like our definition of rival stores, let $G_j = g \neq j \in H : d_{g,j} < \underline{d}^c$ denote the set of contaminated stores near victimized store j , with $G = \bigcup_{j=1}^n G_j$ representing all unique contaminated stores. In this section, we maintain IC.1 and IC.2, but modify IC.3 to include all contaminated stores:

$$\text{IC.3: } i \notin D_{j-}, \text{ where } D_{j-} = \{h \in G : (t_h + w > t_j) \vee (t_h - w < t_j)\}$$

Figure A.9 shows our treatment effect estimates for victimized (Panel A) and rival stores (Panel B) using the revised IC.3 condition. The black lines represent estimates under the condition that contamination status is assigned based on a store’s first nearby crime incident, as with rival treatment. Consequently, each contaminated store is assigned to only one treatment cohort, allowing multiply-contaminated stores to serve as clean controls for subsequent crimes, even if their event windows overlap with store j . Our results show minimal differences from the main specification but exhibit slightly larger standard errors. If we consider contamination from any nearby crime incident with overlapping event windows, the standard errors increase further due to fewer clean controls, yet the effect sizes remain similar (grey lines in Figure A.9). Overall, this reinforces that our main findings are not substantially biased by spillover effects from non-rival stores located within 30 miles of other treated stores.

Figure A.9: Price effects with more restrictive clean control criteria



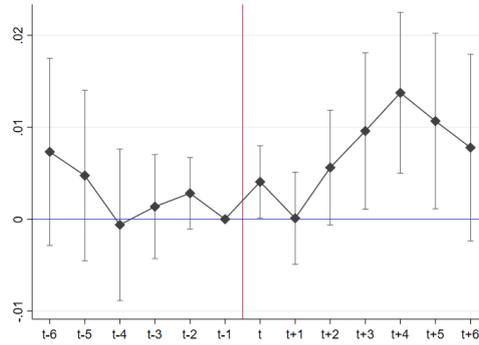
Notes: Panels A and B show cumulative treatment effects, (E_L) , L months after a crime for victimized stores and for rival stores, respectively, along with corresponding 90% confidence intervals based on standard errors clustered at the store level. The black lines show estimates when contamination status is based on a store’s first nearby crime incident. The grey lines show estimates when contamination status is based on any nearby crime incident.

F.3 Variation in rival treatment intensity

Some stores experience multiple nearby crime incidents within a 5-mile radius in a single month. Such underlying differences in treatment intensity violate the SUTVA and may lead

to biased estimates in our rival specification. Therefore, in this section we estimate a variant of our rival specification that excludes rival stores with more than one nearby crime in a given month. The number of rival stores falls from 236 to 219 and four sub-experiments drop out of the estimation sample. However, Figure A.10 shows that the treatment effects, while more volatile, are similar to those from our main specification in Figure 3. Four months after a retail crime incident, prices at rival stores are 1.4% higher than the month before the incident, and the estimate is significant at the 1% level. This suggests that underlying differences in rival treatment intensity are not an important source of bias in our main analysis.

Figure A.10: Rival price effects without multiple contemporaneous rival treatments



Notes: The figure shows the cumulative effect of crime incidents on rival store prices, (E_L) , L months after a crime, along with corresponding 90% confidence intervals based on standard errors clustered at the store level. Compared to our main specification, we exclude rival stores with more than one nearby crime incident in a given month.

G Alternative estimators

One advantage of the stacked DiD estimator, compared to other estimators that address biases from staggered treatment adoption (e.g. Callaway and Sant’Anna, 2021; Borusyak et al., 2024), is that rules defining clean controls readily extend to geographic criteria. This feature allows us to mitigate biases from spillovers to untreated stores and/or differences between stores, while at the same time ensuring a large number of clean control stores. In contrast, with alternative estimators (including canonical TWFE) restrictions on the control group composition must be imposed for the entire sample period, which considerably decreases the number of valid control observations. However, other estimators may be more efficient and better comparable to other studies.

To assess whether our findings are sensitive to the choice of estimator, we use three alternative approaches: i) the canonical two-way fixed effect Difference-in-Differences estimator (TWFE) prone to biases under staggered treatment adoption; ii) the imputation estimator developed by Borusyak et al. (2024) (BSJ); and iii) the estimator proposed by Callaway and Sant’Anna (2021) (CS). We estimate all models separately for both of our treatment groups (victimized and rival stores). Similar to our main specification, we first estimate a distributed lag model for the same event window (equation 2) and then calculate the cumulative treatment effects using the period before treatment as our reference period.

The underlying dataset for the alternative estimators includes all stores, which means the control group may comprise stores within 30 miles of victimized stores. Positive price effects at these stores due to spillovers (see Section 4.1) would imply that we underestimate the treatment effects in these models. To limit the bias from spillovers, we exclude all rival stores in the victimized store specifications and all victimized stores in the rival store specifications. If we restrict the sample to all stores that are located more than 30 miles from any treated stores, the number of clean control stores would fall from 344 to 78 stores, demonstrating the advantages of the stacked DiD framework in our setting.

For the BSJ estimator, we use the accompanying Stata package "did_imputation" from Borusyak et al. (2024). Standard errors allow for clustering at the store level. We calculate standard errors (detailed in Borusyak et al. (2024)) using the treatment effect averages across treatment cohorts (time since treatment) excluding the own unit.⁷ For the CS estimator, we employ the Stata package "csdid" with clustered, asymptotic standard errors and stabilized inverse probability weighting. For both estimators, we consider all never-treated and not-yet-treated stores as potential control observations.

One notable difference among the estimators is the benchmark against which the average distributed lag coefficients are calculated. The canonical TWFE, similar to our main specification, calculates the distributed lag coefficients relative to the last period before the

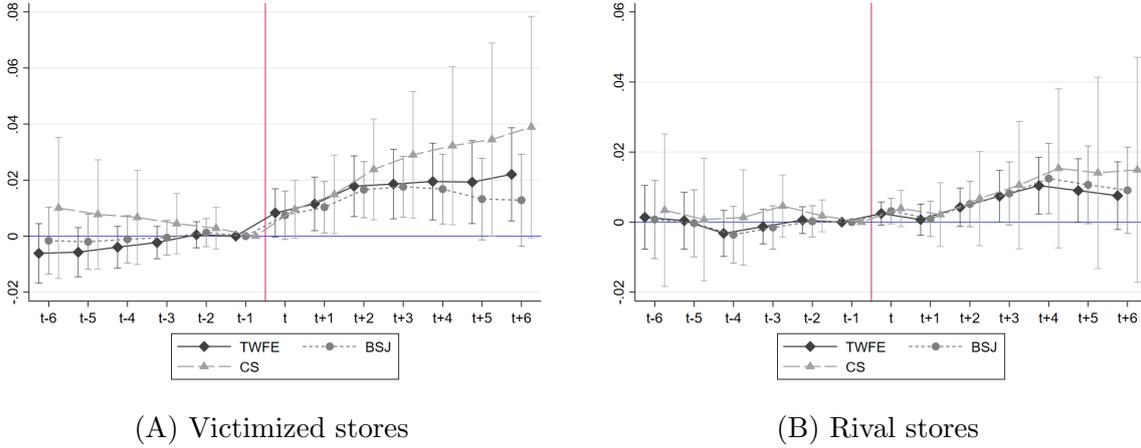
⁷Some observations (one victimized store and a few rival stores) are dropped through the "autosample" option. We cannot impute non-treated potential outcomes for those stores, implying that the BSJ estimator cannot estimate unbiased treatment effects for these stores.

event window. The CS estimator uses the last pre-treatment period as a reference, while BJS bases its comparison on the average of all pre-treatment periods (Roth et al., 2023). If all benchmarks were zero, these differences would be irrelevant. However, since the last pre-treatment period coefficient is slightly negative, the cumulative coefficients from the CS estimator turn out to be higher than those from other estimators. For the CS estimators, we employ "long" differences, yet the interpretation of pre-treatment coefficients still varies slightly between BSJ and other estimators. Importantly, all distributed lag coefficients are insignificant, and pre-tests included in the packages reveal no concerning pre-trends.

Figure A.11 presents the cumulative treatment effects for the alternative estimators. For victimized stores, the effect size four months after the incident is approximately 1.9% using the TWFE estimator and 1.7% with the BSJ estimator (Panel A, Figure A.11). Both coefficients are statistically significant at the 5% level and remain at this elevated level six months after the incident, which aligns with the estimates in our main specification. As expected, we estimate larger effects in the CS specification, with prices approximately 3.2% higher four months after the crime, followed by a slight upward trend over time. Although the standard errors are larger for the CS estimator, its statistical significance is comparable to that of the other estimators.

As anticipated, the treatment effects are slightly smaller than those in our main specification for rival stores. Similar to the main results, effects materialize with some delay, becoming statistically significant at the 5%-level after four months in the TWFE and BSJ specifications, with an effect size of around 1.0-1.2%. Again, estimates are slightly higher in the CS specification but statistically insignificant due to higher standard errors. We find no statistically significant pre-trend coefficients in either the victimized or rival store specifications. Overall, the evidence from the alternative estimators supports our conclusion that the effects of crime on prices are robust and not dependent on the choice of estimator.

Figure A.11: Effect of crime on stores price levels using alternative estimators



Notes: This figure displays the cumulative treatment effects (E_L) L months after a crime on store price levels, using different alternative estimators. The effects are shown with corresponding 90% confidence intervals based on standard errors, which allow for clustering at the store level. Coefficients are interpretable as percentage increases in outcome levels relative to the month before a crime incident. Panel A presents results for victimized stores, while Panel B focuses on rival stores. The black line (TWFE) represents cumulative effects estimated with the canonical two-way fixed effects estimator. The dotted line (BSJ) shows estimates from the imputation estimator proposed by Borusyak et al. (2024), and the light grey line (CS) uses the estimator developed by Callaway and Sant’Anna (2021).

H Additional robustness checks

H.1 Placebo tests

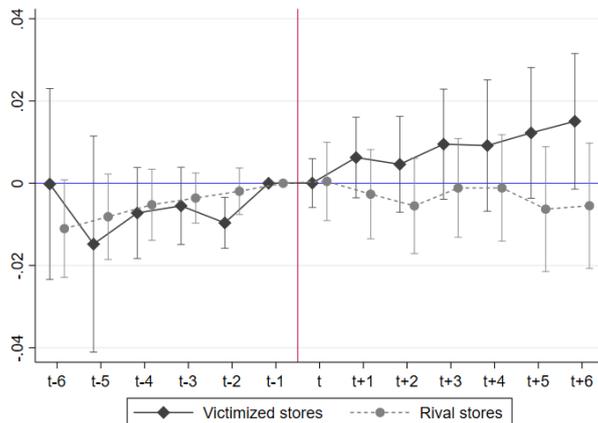
We conduct placebo tests to check for a potential violation of the parallel trends assumption. We shift the treatment assignment forward by 12 months and estimate our main empirical equation (equation 2) with leads and lags based on this placebo treatment assignment. If the placebo treatment timing were to produce large and statistically significant treatment effects, then this would cast doubt on our parallel trends assumption.

We report our placebo test estimates in Figure A.12. The figure indicates that when the treatment timing is moved forward by 12 months, the treatment effects on both victimized and rival stores are negligible and not statistically significant. Notably, victimized stores exhibit a minor positive effect at higher lags, whereas rival stores show a slight negative post-trend, but all coefficients remain statistically insignificant at the 10% level. Overall, the placebo test results support the validity of the parallel trends assumption.

H.2 Longer event window

Figure A.13 illustrates that our main results carry over when we extend the event window from 13 months to 19 months. Prices at victimized stores rise for three months following a crime incident and maintain a higher level for nine months. Rival stores experience similar price increases, albeit with a two-month delay. The positive treatment effects remain intact

Figure A.12: Placebo tests



Notes: The figure shows the cumulative effect of crime incidents on prices, (E_L) , L months after a crime, along with corresponding 90% confidence intervals based on standard errors clustered at the store level. Compared to our main specification, we assign a placebo treatment 12 months prior to a store’s actual treatment date. The black line depicts the cumulative effects of crime on prices at victimized stores, while the grey line represents rival stores.

beyond the original 13-month event window, reinforcing our empirical strategy’s assumption of constant treatment effects outside this period.

The effects on quantity sold (Panel B) are more volatile than those for price, reflecting the much larger variance for the quantity indexes compared to the other indexes. Pre-treatment, the effects at both victimized and rival stores are insignificant and show no significant differences. Interestingly, several months post-crime, we notice a decrease in quantities sold at victimized stores. This reduction is likely due to consumers shifting away from these stores in response to the higher prices. Initially, demand at these stores appears inelastic, but as consumers adjust to the increased prices, demand decreases as expected.

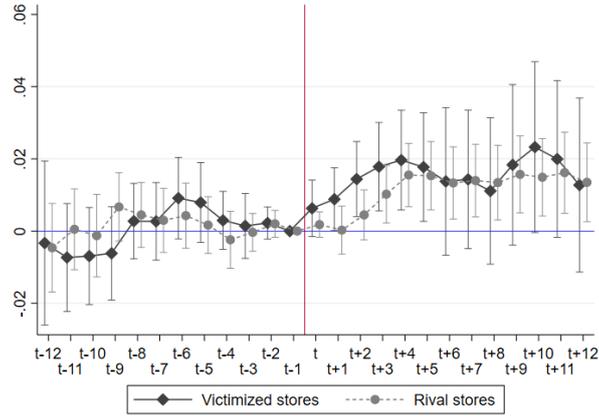
Panel C shows a small (1%) increase in wholesale cost at victimized stores, though the effect is not statistically significant. We view this as unproblematic for two reasons. First, the wholesale cost increase appears two months after the retail price increase, making it unlikely that changes in wholesale cost drive the estimated price effects. Second, the wholesale cost increase disappears by $t + 10$ while the price effect remains stable. Changes in wholesale costs at rival stores are close to zero and not statistically significant.

Taken together, the results from Figure A.13 illustrate that treatment effects materialize and stabilize well before $t + 5$, which validates our choice of an 11-month event window for our main analysis in Section 5.

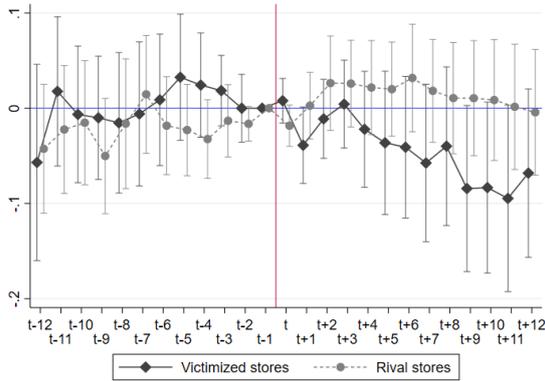
H.3 Sensitivity to outliers

To address potential biases from outliers, we re-estimate our main empirical specification using outcome variables winsorized at the 0.5% and 99.5% levels. We report the results in

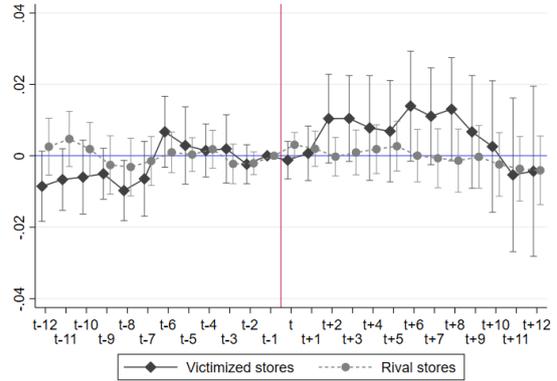
Figure A.13: Effect of crime on store outcomes with extended event window



(A) Price level



(B) Quantity sold

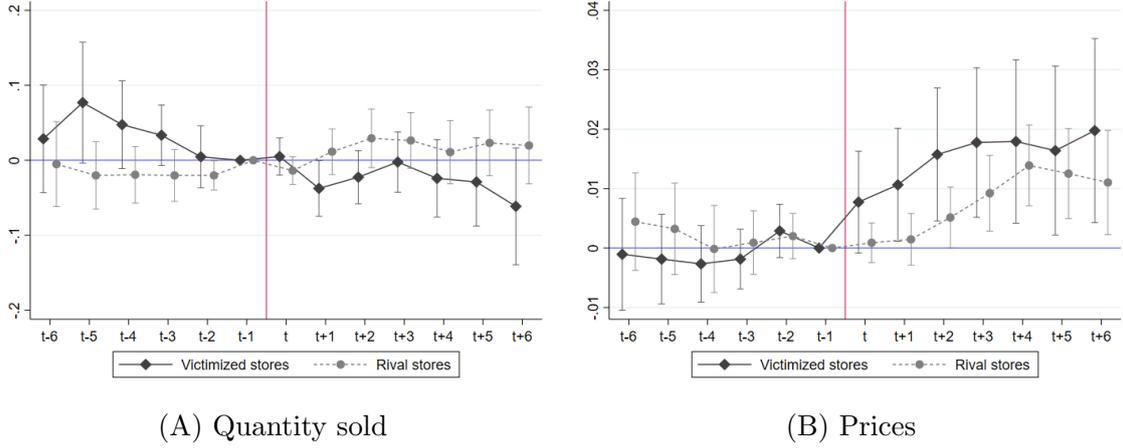


(C) Wholesale cost

Notes: Each panel shows the cumulative treatment effects (E_L) L months after a crime on different outcomes, along with corresponding 90% confidence intervals based on standard errors clustered at the store level. Compared to the main specification, we extend the event window from 11 to 19 months. Coefficients are interpretable as percentage increases in outcome levels relative to the month before a crime incident. The black line depicts the cumulative effects of crime on outcomes at victimized stores, while the grey line represents rival stores. The dependent variables are: store-level price index (Panel A), store-level quantity index (Panel B), and store-level wholesale cost index (Panel C).

Figure A.14. Panel A, corresponding to the estimates in Column 4 of Table 4, indicates that winsorizing the store-level price indexes does not influence the estimated price effects of retail crime, further reaffirming our main results. The results further show that winsorizing decreases the standard errors of our rival treatment estimates, particularly for the quantity index, which has a considerably higher standard deviation than the price and wholesale cost indexes.

Figure A.14: Effect of a crime incident on prices and quantity sold, winsorized



Notes: The figures show cumulative treatment effects, (E_L) , L months after a crime for victimized stores and for rival stores, respectively, along with corresponding 90% confidence intervals based on standard errors clustered at the store level. Panel A shows the effect on quantity sold when the quantity index is winsorized by 1%. Panel B shows the effect on prices when the price index is winsorized by 1%.

H.4 Store closures

One concern is that the costs associated with retail crime incidents may cause victimized stores to go out of business. If the propensity to drop out of the market following a crime is correlated with unobserved store characteristics (e.g. profitability), then our treatment effect estimates may be biased. We view such attrition as unproblematic for two reasons. First, store FE control for time-invariant factors that influence a store’s likelihood of dropping out of the sample. Second, we find that victimized stores are not more likely to drop out of the market compared to non-victimized stores. Of the 56 victimized stores in our main specification, six drop out before the end of the 46-month sample period, a dropout rate of 0.11. For these six stores, the average duration between a crime incident and dropping out is 16 months. Only one store drops out within the first 12 months after an incident, which is equal to 1.8% of victimized stores. 58 of the 311 non-victimized stores drop out during the sample period, a rate of 0.19, or 5% on average per year.⁸ For comparison, the annual dropout rate for restaurants is estimated to be about 30% (Parsa et al., 2005). Taken

⁸The small number of victimized stores dropping out of the sample makes a DiD specification—similar to our main approach—uninformative for testing differential closure rates between treated and control stores.

together, we view this as suggestive evidence that crime incidents do not lead to a higher dropout rate for victimized stores.

H.5 Clustering standard errors at the sub-experiment-by-treatment group level

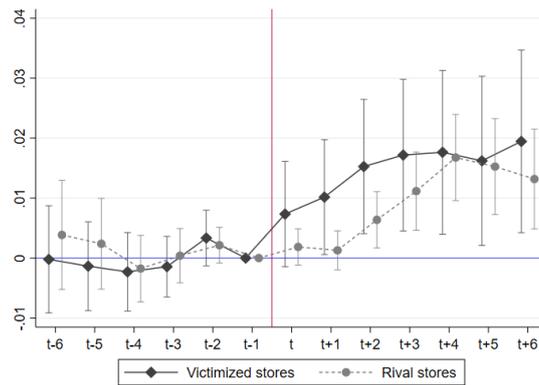
As in other stacked DiD designs, stores in our setting may appear in multiple sub-experiments. A control store in one sub-experiment may serve as a control or rival store in other sub-experiments, and a store treated in one sub-experiment may also serve as a control in another, provided the event windows do not overlap (see Appendix D for details). These “duplicate” stores introduce dependence across sub-experiments, which justifies clustering standard errors at the store level, as discussed by Wing et al. (2024). An alternative approach is to cluster at the sub-experiment-by-treatment group level, as done, for instance, by Cengiz et al. (2019). Both methods perform well in Monte Carlo simulations, provided the number of clusters is sufficiently large (Wing et al., 2024). To account for potential within-group dependence and to assess the robustness of clustering at the store level, we also estimate standard errors clustered at the sub-experiment-by-treatment group level. The resulting standard errors, shown in Figure A.15 are similar to those in our main specification and do not affect our conclusions.

H.6 Grams sold and price per gram

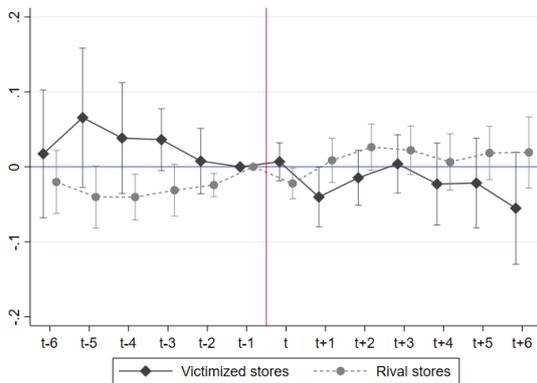
As discussed in Section 3, price and quantity indexes carry numerous advantages over unweighted aggregations and are therefore commonplace in the literature (Leung, 2021; Renkin et al., 2022). Nevertheless, it is important to confirm that our main results are not an artifact of index aggregations. In this subsection, we address this point by using an alternative measure of prices and quantities for the Usable Marijuana product category. Usable Marijuana has the benefit that all products are measured in grams, and hence allow for comparison in units sold in a way that would not be possible in other product categories. We proceed as follows. First, we replace the store-level price index with the store-level (unweighted) mean monthly price per gram of Usable Marijuana. Second, we define quantity sold as the total grams sold of Usable Marijuana per month for each store. Figure A.16 depicts the results when we estimate our main specification using these alternate dependent variables. Panel A shows a similar effect on prices as in the main specification: prices at victimized stores increase by approximately 2% the month of a crime incident and stay at this higher level, while at rival stores prices increase with a delay. Panel B shows no statistically significant or systematic change in grams sold at either victimized or rival stores.

While these findings reinforce our main results, they should be interpreted with caution due to several limitations that contribute to larger standard errors. First, unlike the indexes used in the main analysis, the dependent variables are not weighted by a product’s relative

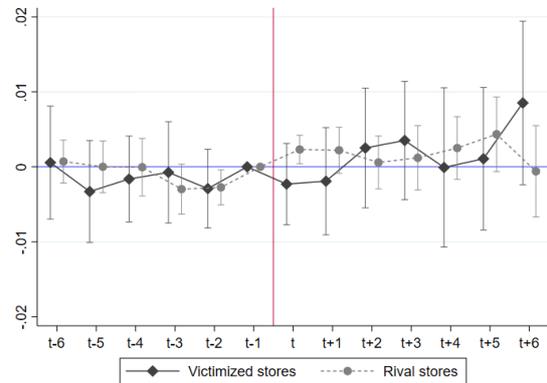
Figure A.15: Effect of crime on rival store outcomes with standard errors clustered at the sub-experiment-by-treatment group level



(A) Price level



(B) Quantity sold

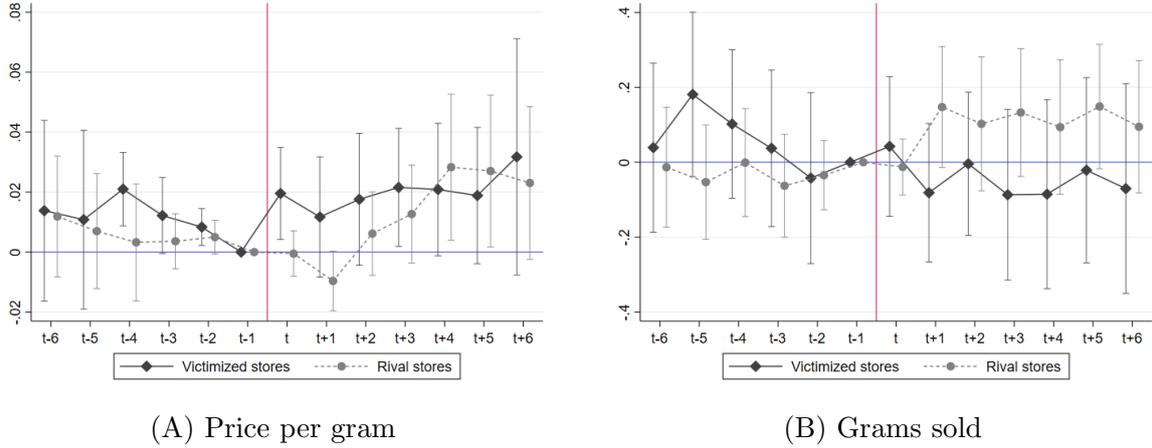


(C) Wholesale cost

Notes: Each panel shows the cumulative treatment effects (E_L) L months after a crime on different outcomes at victimized and rival stores, along with corresponding 90% confidence intervals based on standard errors clustered at the group-by-sub-experiment level. Coefficients are interpretable as percentage increases in outcome levels relative to the month before a crime incident. The dependent variables are: store-level price index (Panel A), store-level quantity index (Panel B), and store-level wholesale cost index (Panel C).

importance to a given store. Thus, a store may experience an increase in grams sold of a product that contributes relatively little to the store’s annual revenue and hence does not factor into major pricing decisions. Second, grams sold and prices per gram exhibit substantial month-to-month variation at the store level, making them noisy measures of quantity sold and prices. Finally, since the Usable Marijuana product category makes up 52% of the cannabis market share (see Table 1), it is difficult to rule out price and quantity changes for other important product categories.

Figure A.16: Effect of crime on price per gram and grams sold



Notes: Each panel shows the cumulative treatment effects (E_L) L months after a crime on different outcomes, along with corresponding 90% confidence intervals based on standard errors clustered at the store level. Coefficients are interpretable as percentage increases in outcome levels relative to the month before a crime incident. The black line depicts the cumulative effects of crime on outcomes at victimized stores, while the grey line represents rival stores. The dependent variables are: store-level monthly average price per gram (Panel A), store-level monthly grams sold (Panel B).

H.7 Price effects of additional crime incidents

Some stores experience more than one crime incident during the sample period (three stores have three subsequent crimes while eight stores have one subsequent crime). In our main analysis, we drop those stores from the sample the month before the second crime incident so as to isolate the effect of the first retail crime on store-level outcomes. In this section, we estimate the effect of subsequent (i.e. non-first) crimes on prices at victimized stores. We drop all stores with a single crime incident from the sample. For the remaining victimized stores, we assign treatment the month of any subsequent crime but we do not assign treatment for the first crime. For some stores, the event window of a subsequent crime overlaps with that of the first crime. Nevertheless, we include observations for the first crime even if the event windows for the first and subsequent crimes overlap. While this introduces potential bias (e.g. due to treatment effects from the first crime that are still evolving when a subsequent crime occurs), dropping observations corresponding to the first crime would reduce the treatment group size to such an extent that inference becomes difficult. We estimate our main specification using

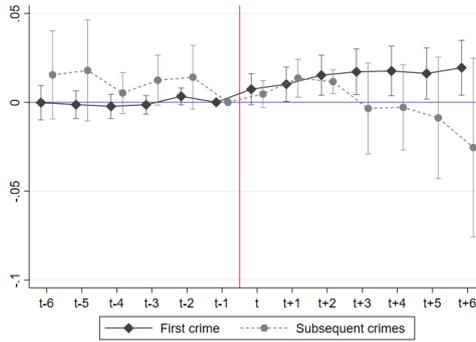
the definition of clean controls described in Section 4.

In Figure A.17 we report the results and compare them to our main estimates from Section 5. The figure shows that from period t through $t+2$ subsequent crime incidents induce a price increase that is statistically significant and similar in magnitude as the first crime. However, the price effects disappear by $t+3$ and become negative (though not significantly different from zero) by $t+6$.

We interpret Figure A.17 as providing suggestive evidence that subsequent crimes may induce additional pass-through. However, these results should be interpreted with caution. Estimates may be biased due to heterogeneous treatment effects, i.e. treatment effects from the first crime incident may still be evolving. Moreover, the subsequent crime treatment group is small, making inference difficult.

Nevertheless, we find it conceptually plausible that subsequent crimes could induce additional costs for retail cannabis stores. These costs can arise if store owners only partially increase security measures following the first crime. Alternatively, criminals may respond to stores' security measures by updating their strategies for theft, which in turn may require additional security costs for retailers. Our crime data provides some evidence for the latter channel. Between 2017 and 2022, armed robberies made up 80 percent of all retail crimes at cannabis stores (see Section 3). Beginning in late 2022, however, there was a gradual increase in the number smash-and-grab incidents. By 2023, smash-and-grab incidents became the dominant form of retail crime at cannabis stores, making up 68 percent of all crimes. Smash-and-grab crimes remained at the same high level in 2024 as well. The changing character of retail crime at cannabis stores suggests that criminals have adjusted their strategies over time, potentially in response to security measures taken at cannabis stores. This, in turn, may induce additional security costs since precautionary security measures for smash-and-grab crimes (e.g. installing concrete barriers) may differ from those for armed robberies (e.g. hiring security guards) (Dowling, 2025). Moreover, smash-and-grab crimes often cause more extensive damage which further increases the costs of such subsequent crimes (see e.g. Saldanha, 2025; Didion, 2025). To the extent that retail crime resembles a dynamic game between criminals and store owners, we view it as plausible that subsequent retail crime incidents induce additional costs for retail stores in other retail sectors as well.

Figure A.17: Price effects of first versus subsequent crime incidents at victimized stores



Notes: The figure shows the cumulative effect of crime incidents on victimized store prices, (E_L) , L months after a crime, along with corresponding 90% confidence intervals based on standard errors clustered at the store level. The black line shows estimates from the first crime incident at a given store. The grey line shows estimates from all subsequent (i.e. non-first) crimes.

I Details of policy analysis—Retail crime as a hidden tax

In this appendix, we provide the details of our welfare analysis. Section I.1 outlines the model used to derive the welfare implications of retail crime pass-through. Section I.2 describes our approach to estimating marginal cost pass-through and presents the corresponding results.

I.1 A model of a hidden crime tax

We consider a market with N firms. Firm j maximizes profits by setting a unidimensional strategic variable, r_j , that can be price, p_j , or quantity, q_j . Each firm produces a single good with marginal costs equal to $mc_j = c'(q_j) + \tau$, where $c'(q_j)$ is the first derivative of the cost function, $c(q_j)$, which is identical for all firms. τ is the unit cost shock stemming from a retail crime incident, that is, the hidden crime tax. We make the simplifying assumption that the hidden crime tax applies equally to all N firms in the affected local market (i.e. stores within a 5-mile radius of the crime).⁹ In this case, Weyl and Fabinger (2013) show that under symmetric imperfect competition the pass-through rate (in dollars) for a small unit tax is:

$$\rho = \frac{dp}{d\tau} = \frac{1}{1 + \frac{\epsilon_D - \theta}{\epsilon_S} + \frac{\theta}{\epsilon_{ms}} + \frac{\theta}{\epsilon_\theta}}, \quad (11)$$

where ϵ_S is the elasticity of the supply function, i.e., of the inverse marginal cost function, and $\epsilon_D = -\frac{p}{qp'(q)}$ is the elasticity of market demand. θ is a conduct parameter summarizing the degree of market competition and can be understood as the ratio of the markup in a

⁹This assumption is supported by our our main results and our marginal cost pass-through estimates (see Section I.2). The former show similar treatment effects between victimized and rival stores, while the latter rules out strategic complementarity in prices as a main mechanism driving the treatment effects.

market to the (fictional) monopoly markup. Consequently, θ is zero for perfect competition and one for the monopoly case. ϵ_θ is the elasticity of the conduct parameter with respect to quantity. ϵ_{ms} is the elasticity of the inverse marginal surplus function equal to $\epsilon_{ms} = \frac{ms}{ms'q}$, which describes the curvature of the demand function. The marginal effect of the unit tax on consumer surplus (CS) and producer surplus (PS) is given by:

$$\frac{dCS}{d\tau} = -\rho q \quad (12)$$

$$\frac{dPS}{d\tau} = -[1 - \rho(1 - \theta)] q \quad (13)$$

Accordingly, the incidence of the unit tax, that is the ratio of consumer to producer harm from an infinitesimal unit tax increase, is given by:

$$I = \frac{\rho}{1 - \rho(1 - \theta)} \quad (14)$$

These results offer important insights regarding the welfare implications of a hidden crime tax. Specifically, they suggest that the pass-through of retail crime is influenced by four factors: the demand elasticity, the supply elasticity, the curvature of the demand function, and the conduct parameter (together with its elasticity with respect to q which is often zero in common models, such as the Cournot model). Under perfect competition ($\theta = 0$), ρ is only determined by the ratio of supply and demand elasticities, where the more inelastic side of the market bears more of the tax burden—a familiar result of the tax literature.

Even if we assume that the crime-induced increase in marginal costs is fully redistributed to other players in the economy (e.g., security service providers) or raises retailers' profits in the case of subjective marginal cost increases, crime-induced price hikes still lead to an excess burden in the case of imperfect competition. This loss results from price distortions caused by the market power of firms. To see this, consider the monopoly case where $\theta = 1$. In this case, the monopolist fully pays the hidden tax out of its profits ($\frac{dPS}{d\tau} = -q$). Yet, consumers still bear $\frac{dCS}{d\tau} = -\rho q$, implying that the tax is more than fully shared by market participants. This excess burden is zero for the perfect competition case, $\theta \rightarrow 0$.

The results of the model also show that the pass-through rate serves as a sufficient statistic (together with θ) for deriving the welfare effects of a unit tax, its incidence, and the excess burden. This has the advantage that one need not impose restrictive assumptions about the underlying market structure.

I.2 Estimating the marginal cost pass-through rate

Next, we estimate the marginal cost pass-through rate which we can directly relate to the sufficient statistic approach from the tax pass-through literature (Weyl and Fabinger, 2013). To estimate the marginal cost pass-through rate, we follow the industrial organization literature

that measures the pass-through of cost shocks and taxes (Ganapati et al., 2020; Hollenbeck and Uetake, 2021; Miller et al., 2017; Muehlegger and Sweeney, 2022).

We estimate the marginal cost pass-through rate, that is, the increase in retail unit prices (in dollars) stemming from a \$1 increase in wholesale unit prices. We specify a model at the store-product-month level that relates a store-product’s retail price to (i) that store-product’s wholesale price and (ii) the average wholesale price paid by rival stores for that same product (with rivals defined as all stores within a 5-mile radius of the focal store). We include rivals’ cost changes to capture the effect of rivals’ cost-induced price changes, i.e. strategic complementarity in prices (Muehlegger and Sweeney, 2022). We estimate the following model for stores in affected markets (i.e. victimized and rival stores):

$$\Delta p_{i,j,t} = \rho \Delta w_{i,j,t} + \beta \Delta w_{i,r(j),t} + \gamma_t + \Delta \varepsilon_{i,j,t}, \quad (15)$$

where $p_{i,j,t}$ is the average price (in dollars) of product i sold at store j in month t , $w_{i,j,t}$ is the average wholesale price that retailer j pays for product i in month t , $w_{i,r(j),t}$ is the average wholesale price that store j ’s rivals pay for product i in month t , and γ_t is the year-month FE.¹⁰ Since the model is in first differences, product and retailer FE are swept out. The pass-through rate, ρ , is the dollar increase in price at store j from a \$1 increase in store j ’s marginal cost. β measures the effect of marginal cost changes at rival stores on store j ’s prices. In Appendix E, we include the costs of stores located further away to gain insight into the geographic scope of sensitivity to costs and to guide our choice of the clean control group discussed in Section 4.

Table A.9 reports the results from the pass-through regression. We find that a \$1 increase in unit wholesale cost corresponds to a retail unit price increase of \$1.67 (Column 1). Such over-shifting of costs onto consumers is in line with the findings of Hollenbeck and Uetake (2021), and indicates that cannabis retailers exercise substantial market power.¹¹ According to equation 11, for pass-through to exceed one it is sufficient that firms have market power ($\theta > 0$), marginal costs are constant, and demand is log-convex ($\epsilon_{ms} < 0$) (Weyl and Fabinger, 2013).

Column 2 shows that the own-cost pass-through estimate is robust to including county-level average wages, the county-level unemployment rate, and the home price index (at the three-digit zip code level) to absorb variation in retail cannabis prices due to local business cycles or changes in house prices (see Stroebel and Vavra, 2019). Similarly, estimates are unchanged when we include region-time FE to account for other spatially correlated shocks that may covary with wholesale and retail cannabis prices (Column 3). We present further

¹⁰Since cannabis transaction data is publicly available, stores have full information on competitors’ unit costs and prices updated on an almost weekly basis. Therefore, we focus on contemporaneous changes in costs and prices. This is in line with pass-through literature from other industries (see e.g. Hollenbeck and Uetake, 2021; Muehlegger and Sweeney, 2022; Conlon and Rao, 2020; Miller et al., 2017).

¹¹Pass-through rates greater than one have been found in a number of empirical studies estimating cost pass-through in other industries (see e.g. Pless and van Benthem, 2019; Conlon and Rao, 2020).

robustness checks in Appendix E, including estimates in levels, in log-log terms, and using store-level price and cost indexes.

We also find economically small but significant effects of rival stores' costs on a store's own retail prices. For a given product sold at store j , a \$1 increase in the unit cost at rival stores corresponds to a \$0.02 increase in store j 's unit price (Table A.9 Column 1). This aligns with Hollenbeck and Uetake (2021) who find that cannabis retailers in Washington behave like local monopolists. Consequently, it is unlikely that the price increase at rival stores after a crime incident reflects a strategic response to increasing prices at victimized stores. Instead, rivals' price increase appears consistent with an own-cost shock e.g. from precautionary security expenditures or higher commercial crime insurance premia. This conclusion remains unchanged across all of the specifications in Table A.9 and Appendix E.

In our welfare analysis, we therefore abstract from strategic complementarity in pricing and assume that the effect of crime on prices at victimized and rival stores runs entirely through the own-cost channel.

Table A.9: Unit cost pass-through rates

	(1) Baseline	(2) Controls	(3) Reg. \times time FE
Own wholesale cost	1.67*** (0.026)	1.67*** (0.026)	1.67*** (0.026)
Rivals' wholesale cost (0-5 miles)	0.016*** (0.006)	0.016*** (0.006)	0.016*** (0.006)
N	5,260,269	5,260,269	5,260,269

Notes: The table reports the estimates of wholesale unit cost pass-through rates from equation 15. We report estimates for own wholesale cost changes and for average changes in wholesale costs at rival stores located within 5 miles of the respective store. All specifications control for month-year fixed effects. Coefficients for Columns 1-4 are interpretable as pass-through rates in dollars. Standard errors are clustered at the store level and shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

I.3 Direct loss from crime and the hidden tax rate

To evaluate whether the magnitude of the estimated hidden crime tax is commensurate with the potential losses from theft if security measures had not been taken, we compare the tax-induced reduction in producer surplus with an approximation of the potential losses from crime.

In our welfare analysis, the fictional annual tax revenue from the hidden crime tax represents increased expenditures on security or insurance. When we divide this measure by the number of affected stores (victimized and rival stores), we estimate that affected stores spend an additional \$42,000 per year on costs associated with crime. The corresponding pass-

through reduces annual producer surplus by \$34,152 per store, a measure that approximates the true cost of security expenditures at the store level.

Next, we compare this measure to the hypothetical loss from crime if security measures had not been taken. We use information on losses from past crime incidents as a lower bound approximation of the costs of hypothetical future crimes. While data on the dollar value of losses or damage from retail crimes are not available, anecdotal evidence from news articles and police reports indicate that most losses from crime incidents at cannabis stores range from \$7,000 to \$25,000 in stolen cash and merchandise (Guzman, 2025; Seattle Police Department, 2021; KOMO News, 2024; Lafferty, 2022) (though in a few rare cases losses have exceeded \$100,000 (Holguin, 2024; Capitol Hill Seattle, 2025)). Under the assumption that losses from future crimes are proportional to those from previous crimes, these losses can be thought of as a conservative measure of the opportunity cost of foregoing security expenditures. It is important to note that the true opportunity cost to stores is typically higher, as it includes other unreported direct and indirect costs, such as physical damage to store infrastructure or psychological stress experienced by employees. Overall, losses from crime cited in news and police reports appear commensurate with the reduction in producer surplus. However, given the lack of hard data on the losses from crime, these results should be interpreted with caution.

I.4 Welfare analysis conditional on victimization risk

Our baseline welfare analysis estimates the average welfare effects across the Washington State retail cannabis sector. Several characteristics of our context, such as its reliance on cash, likely lead to higher crime exposure than in other retail sectors. In this regard, our baseline welfare cost estimate may represent an upper bound.

To better understand how welfare effects vary with victimization risk, we conduct a heterogeneity analysis based on the likelihood that a store is victimized. As shown in Appendix A.3, urban areas face substantially higher risks of victimization. We focus on the divide between King county (Washington’s most populous county) and the rest of the state, classifying stores within King county as high-risk and those outside as low-risk. Although the number of incidents is nearly evenly split between King county (38 incidents) and non-King county areas (37 incidents), the number of cannabis stores outside King County is roughly four times larger. This implies a significantly higher victimization rate within King county—approximately 34.5%—compared to just 9.2% outside.

We exploit this variation to estimate welfare effects for high- and low-victimization risk environments. Following our main specification, we re-estimate treatment effects separately for stores within and outside King County. Four months after a crime, we find a price increase of 2.3% in King County stores and 1.2% elsewhere (see Appendix Figure A.8 for details on the No-King county store specification). This suggests that marginal cost shocks may be larger in high-risk areas, potentially due to a greater need for security investments, larger

insurance premium hikes, or stronger subjective updating of future cost expectations by store owners.

Assuming a constant marginal cost pass-through rate (equal to the average estimated for the full sample) and using average prices in each group (\$29.22 in King County vs. \$27.24 outside), we estimate a hidden “crime tax” of 1.38% for high-risk stores and 0.72% for low-risk stores. Under the assumption that conduct parameters remain constant across locations, and adjusting for market size differences in terms of revenues, this yields a total welfare cost of \$15.68 million and an associated excess burden of \$9.4 million in high-victimization areas. In contrast, low-victimization areas face a welfare cost of \$13.6 million and an excess burden of \$8.1 million.

To generalize to the broader U.S. retail sector, we conduct a simple bounding exercise. If the retail sector more closely resembles low-victimization areas, extrapolating based on the revenue share of non-King County cannabis retailers relative to the national retail sector yields an estimated excess burden of approximately \$43.7 billion—roughly 60% of our baseline estimate of \$73.6 billion. Conversely, if crime effects in other retail sectors are closer in magnitude to those observed in high-victimization areas like King County, the implied nationwide excess burden would rise to \$133.1 billion, or 1.8 times our baseline estimate.

As a further refinement, we divide the broader retail sector into “low-risk” and “high-risk” categories using rankings from Homeland Security Investigations (2022). According to this source, high-risk retail sectors include pharmacy, big box, home improvement, grocery, soft lines, electronics, luxury, specialty, dollar stores, and cell phone stores. We classify all other retail sectors as low-risk sectors, including, e.g., furniture stores, motor vehicle and parts dealers, health and personal care stores (excluding pharmacies), and more. While this classification does not provide exact victimization rates, it allows for an approximate allocation. Using data from the 2020 Annual Retail Trade Survey (U.S. Census Bureau, 2022), we then calculate the annual sales for the high-risk (\$2,524 billion) and low-risk sectors (\$2,006 billion). Applying the low-victimization tax rate to low-risk sectors, the high-victimization rate to high-risk sectors, and the average rate to the remaining sectors yields an intermediate excess burden estimate of \$89.5 billion due to retail crime pass-through.

References

- Axios (2024). Americans’ average daily travel distance, mapped. <https://www.axios.com/2024/03/24/average-commute-distance-us-map>. Retrieved April 24, 2024.
- Borusyak, K., Jaravel, X., and Spiess, J. (2024). Revisiting event study designs: Robust and efficient estimation. *Review of Economic Studies*, pages 1–33.
- Callaway, B. and Sant’Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225:200–230.

- Capitol Hill Seattle (2025). Capitol hill uncle ike's hit in another seattle smash and grab pot shop burglary. *Capitol Hill Seattle Blog*. <https://www.capitolhillseattle.com/2025/05/capitol-hill-uncle-ikes-hit-in-another-seattle-smash-and-grab-pot-shop-burglary/>.
- Cengiz, D., Dube, A., Lindner, A., and Zipperer, B. (2019). The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics*, 134:1405–1454.
- Conlon, C. T. and Rao, N. L. (2020). Discrete prices and the incidence and efficiency of excise taxes. *American Economic Journal: Economic Policy*, 12:111–143.
- Didion, A. (2025). Two seattle pot shops hit in crash-and-grab burglary attempts within hours. *KING 5 News*. <https://www.king5.com/article/news/crime/two-seattle-pot-shops-crash-grab-burglary-attempts/281-16d9efaf-b95d-4c1f-a867-7d71c28fd53c> (Accessed 12 June 2025).
- Dowling, J. (2025). Thieves push cement barriers through maple valley pot shop to get inside. *FOX13 Seattle*. <https://www.fox13seattle.com/news/thieves-push-barriers-pot-shop> (Accessed 12 June 2025).
- Ganapati, S., Shapiro, J. S., and Walker, R. (2020). Energy cost pass-through in US manufacturing: Estimates and implications for carbon taxes. *American Economic Journal: Applied Economics*, 12:303–342.
- Guzman, A. (2025). Smash-and-grab robbery damages kemp's cannabis in seattle's belltown neighborhood. *FOX 13 Seattle*. <https://www.fox13seattle.com/news/smash-and-grab-kemps-cannabis> (Accessed 12 June 2025).
- Holguin, B. (2024). Seattle pot shop owner faces \$100,000 loss after thieves smash in to take atm. *KIRO 7 News*. <https://www.kiro7.com/news/local/seattle-pot-shop-owner-faces-100000-loss-after-thieves-smash-take-atm/VNFFMA3KVJFE7GXB677SBI76RE/> (Accessed 12 June 2025).
- Hollenbeck, B. and Uetake, K. (2021). Taxation and market power in the legal marijuana industry. *RAND Journal of Economics*, 52:559–595.
- KOMO News (2024). Sodo pot shop targeted by crash-and-grab burglars for 2nd time in 5 months. *KOMO News*. <https://komonews.com/news/local/seattle-sodo-pot-shop-burglary-crash-and-grab-crime-crisis-stolen-vehicle-suspects-c>
- Lafferty, K. (2022). Washington pot shops call on the state to improve protection from violent robberies. *King 5 News*. <https://www.king5.com/article/news/crime/washington-pot-shops-call-on-the-state-to-improve-protection-from-violent-robberies/281-20fc84ad-d287-4990-bc1b-7ca96875d6f1>.

- Leung, J. H. (2021). Minimum wage and real wage inequality: Evidence from pass-through to retail prices. *Review of Economics and Statistics*, 103:1–16.
- Lopez, E., Boxerman, R., and Cundiff, K. (2023). Shoplifting trends: What you need to know. *Council on Criminal Justice*.
- Miller, N. H., Osborne, M., and Sheu, G. (2017). Pass-through in a concentrated industry: empirical evidence and regulatory implications. *RAND Journal of Economics*, 48:69–93.
- Muehlegger, E. and Sweeney, R. L. (2022). Pass-through of own and rival cost shocks: Evidence from the us fracking boom. *Review of Economics and Statistics*, 104(6):1361–1369.
- National Retail Federation (2024). The impact of retail theft and violence. https://297051953189d612da9e-1e2a7931911c2abaf913026fb7c64860.ssl.cf1.rackcdn.com/Research/Retail%20Theft%20%26%20Violence/NRF_ImpactofRetailTheftViolence_2024.pdf (Accessed 4 June 2025).
- Parsa, H. G., Self, J. T., Njite, D., and King, T. (2005). Why restaurants fail. *Cornell Hotel and Restaurant Administration Quarterly*, 46:304–322.
- Pless, J. and van Benthem, A. A. (2019). Pass-through as a test for market power: An application to solar subsidies. *American Economic Journal: Applied Economics*, 11(4):367–401.
- Renkin, T., Montialoux, C., and Siegenthaler, M. (2022). The pass-through of minimum wages into U.S. retail prices: Evidence from supermarket scanner data. *Review of Economics and Statistics*, 104:890–908.
- Roth, J., Sant’Anna, P. H., Bilinski, A., and Poe, J. (2023). What’s trending in difference-in-differences? A synthesis of the recent econometrics literature. *Journal of Econometrics*, 235(2):2218–2244.
- Saldanha, A. (2025). Whatcom county marijuana shop burglaries show distinct pattern; suspects remain at large. <https://www.bellinghamherald.com/news/local/crime/article298197908.html> (Accessed 9 June 2023).
- Seattle Police Department (2021). Detectives need your help identifying lake city pot shop robbery suspects. SPD Blotter. <https://spdblotter.seattle.gov/2021/11/23/detectives-need-your-help-identifying-lake-city-pot-shop-robbery-suspects/> (Accessed 23 June 2025).
- Stroebel, J. and Vavra, J. (2019). House prices, local demand, and retail prices. *Journal of Political Economy*, 127:1391–1436.

- U.S. Census Bureau (2022). 2020 Annual Retail Trade Survey (ARTS) Data. Accessed July 2025. Includes retail sales, expenses, inventories, and e-commerce data for 2020.
- U.S. Department of Homeland Security (2022). Detecting and Reporting the Illicit Financial Flows Tied to Organized Theft Groups (OTG) and Organized Retail Crime (ORC) [Report]. <https://www.acams.org/en/media/document/29436> (Accessed 10 April 2024).
- Washington State Department of Health (2024). Tobacco and cannabis use dashboard. <https://doh.wa.gov/data-and-statistical-reports/washington-tracking-network-wtn/tobacco-and-cannabis/dashboard> (Accessed 24 April 2024).
- Washington State Office of Financial Management (2024). Population of cities, towns and counties. https://ofm.wa.gov/sites/default/files/public/dataresearch/pop/april1/ofm_april1_population_final.pdf.
- Weyl, E. G. and Fabinger, M. (2013). Pass-through as an economic tool: Principles of incidence under imperfect competition. *Journal of Political Economy*, 121(3):528–583.
- Wing, C., Freedman, S. M., and Hollingsworth, A. (2024). Stacked difference-in-differences. *NBER Working Paper w32054*.