

The Pass-through of Retail Crime*

Carl Hase[†] Johannes Kasinger[‡]

November 2024

Abstract

This paper shows that retailers increase prices in response to organized retail crime. We match store-level crime data to scanner data from the universe of transactions for cannabis retailers in Washington state. Using quasi-experimental variation from robberies and burglaries, we find a 1.5-1.8% price increase at victimized stores and nearby competitors. This rise is not driven by short-to-medium-term demand changes but is consistent with an own-cost shock. Effects are larger for independent stores and less concentrated markets. We estimate that crime imposes a 1% ‘hidden’ unit tax on affected stores, implying \$28.4 million additional social costs, primarily borne by consumers.

Keywords: Organized retail crime, public crime prevention, social costs of crime, pricing, market power, tax incidence

JEL Codes: D22, H22, H32, H76, L11, L81, K42

*We thank Thorsten Schank, Thomas Otter, Guido Friebel, Markus Reisinger, Michael Pollmann, Frank Verboven, Sylvia Hristakeva, Aljoscha Janssen, Bart Bronnenberg and seminar participants at Tilburg University, UW Foster School of Business, NASMES, CAED, ESEM, EARIE, EMAC, and Marketing Science for helpful comments. We are grateful to Tim Haggerty, Chuck Groom, Erik Skaar, and Ian Eisenberg for invaluable insights on the cannabis industry and the traceability data.

[†]Goethe University Frankfurt and JGU Mainz, hase@uni-frankfurt.de

[‡]Tilburg School of Economics and Management - Tilburg University, j.kasinger@tilburguniversity.edu

1 Introduction

Retail crime has surged to the forefront of public discourse in the United States, fueled by an apparent rise in organized retail crime.¹ Retail crime imposes costs on businesses, individuals, and society. A comprehensive assessment of these costs is central for determining the optimal level of public crime prevention (Stigler, 1970; Owens and Ba, 2021). One factor often overlooked when considering the social costs of retail crime is its impact on market outcomes. In particular, crime-induced price changes, i.e., the pass-through of retail crime, have critical distributional implications and can introduce an excess burden by distorting firms’ and consumers’ decisions. Yet, evidence of a causal link between retail crime and market prices is nonexistent, and the associated welfare implications and underlying mechanisms are not well understood.

This study investigates the impact of organized retail crime on prices and the resulting welfare implications using the Washington state retail cannabis industry as a natural laboratory.² This setting overcomes key obstacles that have limited previous research on the effects of retail crime on market outcomes: (1) the ability to observe the time and location of retail crime incidents, and (2) the need for comprehensive information on prices and quantities for the universe of firms. We match store-level data on reported robberies and burglaries with scanner data capturing the universe of transactions for every cannabis store. Such rich data on crime incidents, prices, and quantities—which is not available in any other industry—enables a clean identification of the effects of retail crime on store-level outcomes and their underlying mechanisms. Our setting further allows us to conduct a comprehensive welfare analysis using a sufficient statistics approach based on insights from the tax incidence literature (Weyl and Fabinger, 2013).

Exploiting quasi-experimental variation from the timing of store-level crime incidents, we use a difference-in-differences framework and compare outcomes at victimized stores to outcomes at stores in unaffected local markets.³ We estimate that victimized stores increase prices by 1.8% four months after a crime incident. The price increases persist and remain stable twelve months after an incident. We find no immediate effect on quantities sold in the months following a crime, which speaks against fear-induced demand substitution on the part of the consumers. However, quantities sold fall over a longer time horizon, consistent

¹In contrast to petty theft or shoplifting, organized retail crime involves large-scale theft of merchandise or cash with the intent to resell, often as part of a criminal enterprise. In line with this definition, we limit our analysis to robberies and burglaries. Throughout the paper, we use the terms “organized retail crime” and “retail crime” interchangeably. More details on organized retail crime can be found in Section 2.2.

²The US cannabis industry generates over \$25 billion in annual revenue. More than half of Americans live in states with legal retail cannabis markets, and about one-third of adults regularly consume cannabis in those states (Chapekis and Shah, 2024; Statista, 2024).

³In contrast to other studies exploiting spatial treatment variation, a major advantage in our setting is that we observe the universe of vertical transactions between producers and retailers. This allows us to estimate retailers’ spatial sensitivity to competitors’ unit cost shocks and identify a valid control group that is not contaminated by strategic price competition (Hollenbeck and Uetake, 2021; Muehlegger and Sweeney, 2022).

with consumers responding to higher prices. Additionally, we find a negligible effect of crime incidents on wholesale costs at victimized stores.

We also investigate effects at nearby rival stores, e.g. due to strategic price responses or an own-cost shock. We find that rivals of victimized stores—stores within a 5-mile radius of a victimized store—increase prices by a similar amount (1.5% after four months). However, in contrast to victimized stores, this price increase materializes with a lag of around two months. Again, we find no medium-term effects on wholesale costs or quantities sold. This finding makes demand substitution an unlikely mechanism for driving the price response at rival stores. To assess whether the price response at rival stores is strategic, we leverage the fact that in addition to retail prices and quantities, we also observe the wholesale prices that each retailer pays for its products.⁴ This allows us to estimate the extent to which changes in stores’ own costs and the costs of their competitors are passed through to retail prices. Our results indicate that stores primarily adjust prices in response to their own costs rather than rivals’ costs. We find a cost pass-through rate of 1.67, i.e. a \$1 increase in a retailer’s wholesale unit cost corresponds to a \$1.67 increase in its retail price. The influence of rivals’ cost changes is statistically significant but economically negligible at \$0.02 (from a \$1 increase in wholesale price). Thus, rivals’ substantial price hikes due to nearby retail crimes cannot be attributed to a strategic price response.

We further identify heterogeneity in price effects conditional on store and local market characteristics. First, we find that crime pass-through is higher at independent stores as compared to chains. Our findings suggest that retail crime pass-through is particularly relevant for mom-and-pop shops, while owners of multiple stores can offset costs across locations, mitigating retail crime’s impact on prices. Second, we find the largest price increases for stores operating in markets with comparably low market concentrations. This finding aligns with the idea that pass-through rates and market power are inversely related, as increased competition makes prices more sensitive to marginal costs—a theoretical prediction that relies on the curvature of firms’ cost functions (Ritz, 2024) and on whether pass-through rates exceed unity (Miller et al., 2017).

Our main specification builds on the stacked difference-in-differences (DiD) estimator introduced by Cengiz et al. (2019) with treatment timing defined as the store-specific month of a crime incident. Like other DiD frameworks, stacked DiD estimates a causal treatment effect under the assumption of parallel trends and no anticipation. To account for biases from heterogeneous treatment timing (De Chaisemartin and d’Haultfoeuille, 2020; Goodman-Bacon, 2021), stacked DiD creates crime incident-specific sub-experiments comparing treated stores to “clean” control stores, i.e., never-treated or not-yet-treated stores for a particular crime incident. We choose stacked DiD over related estimators (e.g., Callaway et al., 2024; Borusyak et al., 2024) because the rules for defining clean controls can be readily extended to geographic criteria. In addition to the timing-based criteria, we require incident-specific

⁴The cannabis industry is vertically disintegrated and consists of two types of firms: upstream producer-processors and downstream retailers. See Section 2 for details on the cannabis industry.

clean control stores to be located between 30 to 60 miles from the respective victimized store (which we refer to as unaffected local markets). These conditions account for confounding effects from treatment spillovers to nearby stores (Muehlegger and Sweeney, 2022) while ensuring a sizeable control group that is comparable to our treatment group.⁵ Our paper is one of the first to extend the stacked DiD framework to spatial criteria, contributing to the advancing literature on this methodology (Cengiz et al., 2019; Deshpande and Li, 2019; Butters et al., 2022; Wing et al., 2024).⁶

In our welfare analysis, we postulate that retail crime incidents constitute a positive marginal cost shock to affected stores that can be understood as a hidden crime tax. While we abstract from fixed cost shocks in our analysis, as they do not affect pricing decisions according to standard theory, we remain agnostic about the source of the marginal cost shock. This shock could arise from direct marginal costs, proportional to quantity sold, such as increased spending on security and insurance, as suggested by various industry surveys, reports, and news articles documenting retailers’ efforts to combat crime (U.S. Chamber of Commerce, 2023; Dunham, 2021; National Retail Federation, 2022; U.S. Department of Homeland Security, 2022). However, it may also reflect retailers’ subjective perceptions of their marginal costs that could be influenced by total costs or other factors.

To derive the welfare implications of retail crime pass-through, we draw from the imperfect competition model by Weyl and Fabinger (2013), which nests standard imperfect competition models, as well as the monopoly and perfect competition cases. The implications of standard tax theory carry over to our case: i) the burden of the hidden crime tax falls disproportionately on the more inelastic side of the market; ii) under imperfect competition, there is an excess burden that increases with firms’ market power; iii) the pass-through rate serves as a sufficient statistic for tax incidence and the excess burden.

Next, we employ the model—together with our marginal cost and retail crime pass-through estimates—to quantify the welfare effects of the crime-related price increases in the Washington state cannabis industry. At the average unit price, our estimated retail crime pass-through rate implies a \$0.42 unit price increase in affected markets. This unit price increase corresponds to a hidden unit tax of approximately \$0.25, or about 1% of the average unit price. We estimate an annual decrease in consumer surplus of about \$19.1 million and a negative effect on retailers of around \$9.3 million. These results indicate an incidence of the hidden crime tax on consumers of around 67% and a total social cost of retail crime pass-through of \$28.4 million. The estimated annual excess burden is around \$17.0 million, assuming that the entire fictional tax revenue stays within the Washington state economy.

⁵We conduct several robustness checks that show that our main findings are not sensitive to our choice of estimator or the definition of unaffected local markets.

⁶Using a stacked DiD estimator, Deshpande and Li (2019) investigate spillover effects of social security field office closings to nearby offices. However, they use timing-based clean control inclusion criteria that do not incorporate geospatial distance. Our estimator combines both timing and geospatial distance in defining clean controls. Butters et al. (2022) investigate spillover effects of state excise taxes on chain stores in nearby states. Their approach is similar to ours in that they exclude chain stores whose parent company is affected by a tax hike in another state within the event window.

The implications of our findings extend well beyond the cannabis industry. Organized retail crime is common in many retail sectors and has triggered a strong policy response (see Section 2.2). Numerous industry surveys suggest that retailers across different industries invest heavily in strategies to prevent retail crime and pass these additional costs through to consumers (U.S. Department of Homeland Security, 2022; Dunham, 2021). This conjecture is supported by the findings of a study by Jackson and Tran (2020), which shows that increased felony larceny thresholds are associated with higher prices for automobiles and computers. The similarities between the retail cannabis industry and other retail settings in terms of cost structures and demand elasticities, as well as the absence of a prevalent black market competing with legal cannabis sales, further support the generalizability of our findings (Hollenbeck and Uetake, 2021). A naive extrapolation of our annual excess burden estimate, using relative sales shares and assuming retail crime impacts other retail industries similarly, suggests that the additional U.S.-wide social costs of retail crime pass-through exceed \$68 billion. While effects are likely to be less pronounced in other sectors due to different market structures and retail crime incidence, the estimate highlights that the social costs of retail crime are significantly underappreciated when not accounting for pass-through effects.

We contribute to two broad strands of literature. First, our study contributes to the extensive literature on the economic impact of crime and the related policy discussion on the optimal level of crime prevention. Much of this literature focuses on estimating the social costs of crime and the trade-offs between the costs and benefits of crime prevention (Stigler, 1970; Owens and Ba, 2021). We are the first to document a causal relationship between retail crime and market prices. Additionally, we provide a thorough discussion of the associated welfare implications highlighting an often neglected aspect of retail crime’s social costs.

Previous research has linked property and violent crimes to a range of economic outcomes: increased property prices (Lynch and Rasmussen, 2001; Gibbons, 2004; Linden and Rockoff, 2008), economic growth (Fenizia and Saggio, 2024), urban depopulation (Cullen and Levitt, 1999), elevated savings levels (De Mello and Zilberman, 2008), changes in working time (Hamermesh, 1999), land use and crop yields (Dyer, 2023), and even increased physical activity (Janke et al., 2016). More specific to our focus, several studies examine crimes’ impact on consumer behavior and business dynamics. Mejia and Restrepo (2016) highlight how crime reduces the consumption of conspicuous goods. Fe and Sanfelice (2022) observe a crime-related decline in local food and entertainment consumption, while others identify a connection between crime and overall business activity and entrepreneurial decisions (Greenbaum and Tita, 2004; Hipp et al., 2019; Rosenthal and Ross, 2010). Using increased public security expenditures as an instrument for violent crime in Columbia, Rozo (2018) shows that increased violence leads to lower output and prices, along with a rise in firms exiting the market. Stolkin (2023) links organized drug trafficking to higher consumer prices in Mexico, and Jackson and Tran (2020) find larceny thefts are positively associated with prices for

computers and cars. In contrast to existing studies, we focus on organized retail crime in the U.S. and identify underlying factors behind the rise in retail prices.

Second, we contribute to the extensive literature studying the pass-through of cost shocks, the underlying drivers, and their implications (Weyl and Fabinger, 2013; Miravete et al., 2018; Miller et al., 2017). In our study, we examine a distinctive type of cost shock—retail crime—where costs are endogenous to firms updating their beliefs about the probability of future crime risks and outcomes. Our analysis suggests that rivals’ price adjustments are less a strategic response and more a reaction to their own increased marginal costs (e.g. due to precautionary security or insurance expenditures). This pattern could reflect a kind of spatial knowledge or learning spillover (Thornton and Thompson, 2001; Audretsch and Feldman, 2004; Bloom et al., 2013). Such spillovers are not only relevant in the context of retail crime but also in other settings where firms compete in local markets and learning from competitors is central to firms’ costs. In these cases, price changes driven by spillovers could be mistakenly interpreted as strategic pricing.

Our heterogeneity analyses also contribute to the growing literature on asymmetric strategies and market outcomes between chain and independent stores (e.g. Jia, 2008; Hollenbeck, 2017; Hollenbeck and Giroldo, 2022; Janssen and Zhang, 2023; Klopach, 2024). We find that the pass-through of retail crime is considerably smaller for chain stores, aligning with studies that show a lack of within-chain price adjustments to local conditions (Hitsch et al., 2021; DellaVigna and Gentzkow, 2019). From a welfare perspective, this pattern suggests a potential benefit to increasing the presence of chain stores within an industry. Additionally, our finding of higher pass-through rates in less concentrated markets adds to the ongoing discussion about the relationship between market competition and pass-through rates (Miller et al., 2017; Cabral et al., 2018; Genakos and Pagliero, 2022; Ritz, 2024).

This paper proceeds as follows. Section 2 describes the institutional context of our study, including retail crime in the United States and the retail cannabis industry. Section 3 details our data and Section 4 describes our main empirical strategy. Section 5 presents our main findings. Section 6 discusses potential endogeneity concerns and robustness checks. Section 7 outlines our policy analysis and derives the welfare implications. Section 8 concludes.

2 Institutional Background

2.1 The cannabis industry in Washington state

Approximately 50% of U.S. states have legal recreational cannabis markets. Washington state’s cannabis market opened in July 2014 for adults 21 years and older. Cannabis has since become one of the largest agricultural industries in the state, contributing \$1.85 billion to gross state product (Nadreau et al., 2020). Cannabis consumption is widespread, with approximately 30% of Washington adults consuming cannabis on a monthly basis (Washington State Department of Health, 2024). Consumption is relatively equal across race/ethnicity,

education, and gender, but decreases monotonically with income and age. We detail the demographic characteristics of cannabis consumers in Appendix A.1.

The industry is regulated by the Washington State Liquor and Cannabis Board (LCB) which offers separate business licenses for upstream and downstream establishments (Washington State Legislature, Title 314, Chapter 55). Producer-processors (i.e. upstream establishments) can cultivate, harvest, and process cannabis but cannot sell to end consumers. Retailers (i.e., downstream establishments) can purchase fully packaged and labeled products from producer-processors and sell them in retail stores. Producer-processors cannot own retail licenses and vice versa, creating complete vertical separation along the supply chain. Retailers can only buy from producer-processors located in Washington state and producer-processors can only sell to retailers in the state. This seals off the core of the supply chain from other U.S. states with legal recreational markets. Retailers cannot sell online, meaning retailers operate brick-and-mortar stores. Appendix A.1 contains additional information on the cannabis supply chain.

Cannabis business licenses are capped by the LCB at 556 retailers and 1,426 producer-processors (Washington State Liquor and Cannabis Board, 2020). Licenses are granted at the establishment level so that a single firm can own several licenses. However, a firm can only own licenses of the same type. Approximately 65% of retail stores belong to one- or two-store firms; 25% of stores belong to 3-5 store chains; less than 11% of stores belong to chains with more than 5 stores.⁷ Not all licenses are actively in business, meaning that some license holders have not opened an establishment and have no reported sales activity, especially at the producer-processor level. During our sample period, there were 511 active retailers and 692 active producer-processors.

The LCB distributes retail licenses to counties according to population density but there are no restrictions on where producer-processors can be located. Retailers are located in 37 of the 39 counties in Washington state. The geographic distribution of retail stores is illustrated in Figure 2, Panel A.

Retail sales are subject to a 37% sales tax but there is no tax on upstream sales. Per month, retailers sell approximately 14,500 units and earn about \$280,000 in (tax-inclusive) revenue (see Table 1).⁸ For more information on the distribution of store characteristics, see Appendix A.2.

Table 1 provides an overview of the cannabis product market. Retail stores sell a variety of cannabis products—around 470 distinct products per month on average. The LCB classifies products according 12 categories (Washington State Legislature, 2015). Usable marijuana (i.e. unprocessed dried flower) and concentrate for inhalation (e.g. for use in vape pens) account for more than 80% of all retail sales. Another 14% of retail sales comes from solid

⁷When the market was created in 2014, the LCB allocated licenses according to a lottery. Since a single firm could apply for more than one license, the lottery created exogenous variation in firm size (see e.g. Hollenbeck and Giroldo, 2022).

⁸For context, the average cannabis retailer generates about twice the revenue of an average convenience store or one-fifth of an average supermarket in the United States (Statista, 2022, 2024).

edibles (chocolate bars, gummies, etc), liquid edibles (soda and other infused drinks), and infused mix (e.g. pre-roll joints infused with concentrates). The remaining categories make up less than 2% of total revenue; these are topical products (e.g. creams and ointments), packaged marijuana mix (e.g. pre-roll joints), capsules, tinctures, transdermal patches, sample jar, and suppository.

Table 1: Product market descriptive statistics

(A) Retail store characteristics

	Monthly average per store	Sample total
Establishments		511
Units sold	15,157 (5,829)	263 million
Distinct products	470 (305)	210,842
Sales	\$282,857 (\$273,082)	\$5.07 billion

Notes: Column 1 reports monthly averages at the store level during the sample period. Standard deviations are in parentheses. Column 2 reports totals across all stores and months in the sample period. Sales are tax-inclusive.

(B) Retail product categories

Product category	Monthly sales (in millions of \$)	Market share
Usable marijuana	\$58.77	0.52
Concentrate for inhalation	\$34.70	0.31
Solid edible	\$8.45	0.08
Infused mix	\$5.40	0.05
Liquid edible	\$2.96	0.03
Other	\$2.16	0.02

Notes: Column 1 reports the average monthly retail cannabis sales (in millions of dollars) across Washington state for the major product categories during the sample period; Column 2 shows the corresponding market shares. Sales are tax-inclusive.

2.2 Organized retail crime

Organized retail crime in the United States

Organized retail crime is defined as large-scale theft of retail merchandise or cash, typically by two or more people, often as part of a criminal enterprise. While the exact definition varies

across law enforcement agencies and jurisdictions, a defining characteristic of organized retail crime is that merchandise is not stolen for personal use but is instead resold through third-party outlets. This contrasts with petty theft (e.g. stealing cosmetics because one cannot afford it), which is not the focus of our study. For ease of notation, we refer to *organized retail crime* and *retail crime* interchangeably throughout the paper.

In 2022, the U.S. Chamber of Commerce declared organized retail crime a “national crisis” (U.S. Chamber of Commerce, 2022). According to the National Retail Security Survey (NRSS), organized retail crime is responsible for nearly \$5 billion in annual losses (National Retail Federation, 2020). Organized retail crime is pervasive in scope, with criminals targeting a variety of stores and goods (U.S. Department of Homeland Security, 2022). Incidents are often violent in nature and threaten employee and customer safety.⁹ Many retailers invest heavily in strategies to prevent retail crime, including third-party guard services, enhanced surveillance technologies, locking cases, and employee safety and de-escalation training (see e.g. Target, 2023; Reagan and Schlesinger, 2019; Fonrouge, 2022).¹⁰

Many states have enacted legislation specifically targeting retail crime, including stiffer penalties for people caught stealing from stores with the intent to resell merchandise, and adding language targeting organized groups that rob multiple retail outlets (Lewis, 2023). In addition to legislation, numerous local, state, and federal law enforcement agencies have created retail crime task forces across the United States (Washington State Office of the Attorney General, 2022; U.S. Department of Homeland Security, 2022).

Organized retail crime in cannabis


Similar to other retail sectors, retail crime is pronounced at cannabis stores. Between 2017 and 2023 there were 210 reported robberies and burglaries at cannabis retailers in Washington state. Figure 1 illustrates that retail crime in cannabis has increased over time and that the trend largely tracks overall robberies in Washington state. This trend has prompted policymakers to take action. Retail cannabis was cited by the Washington state attorney general when creating the state’s retail crime task force (Washington State Office of the Attorney General, 2022), and in 2023 a senate bill was introduced in the state legislature aimed at providing tax relief for security improvements at cannabis retail stores (Washington State Legislature, 2023).

Organized retail crimes in cannabis are predominantly armed robberies that occur during store hours. Perpetrators typically extract cash from the register, demand that employees

⁹In 2023, 81% of NRSS respondents reported an increase in retail crime-related violence against employees (National Retail Federation, 2023). In a separate survey, more than three-quarters of retailers stated that a criminal had threatened to use a weapon against an employee, while 40% of Asset Protection Managers reported incidents where an organized retail criminal used a weapon to inflict harm on an employee (Dunham, 2021).

¹⁰In 2023, for example, 34% of National Retail Security Survey respondents increased payroll to support security and 45% increased the use of third-party security personnel as a measure of crime prevention (National Retail Federation, 2023).

Figure 1: Comparing overall robberies to cannabis retail crime in Washington state



Notes: The figure displays annually reported organized retail crime incidents in the Washington state cannabis industry (solid line) and robberies per 100,000 inhabitants in Washington state (dotted line) from 2017 through 2022. Data sources: Uncle Ike’s Robbery Tracker for cannabis industry incidents and the FBI Crime Data Explorer for state-wide robbery statistics.

give them access to the safe, or clear shelves of merchandise. Since cannabis remains illegal at the federal level, cannabis stores do not have access to electronic payment processing services like Visa and Mastercard and nearly all retail transactions are cash-based.¹¹ Stores thus handle a large amount of cash, making them prime targets for retail crime.

Robberies are often violent: in several instances, consumers or employees have been temporarily taken hostage, physically assaulted, shot at, or even killed. Victimized stores often respond to retail crime incidents by increasing security expenditures. According to the state’s Organized Retail Crime Task Force, interviews with store owners, and numerous news reports, the most common expenditure involves hiring additional security guards (Washington State Liquor and Cannabis Board, 2022; Saldanha, 2022; Phan, 2022). Since contracted security guards often cost between \$75-100 per hour in Washington state, this amounts to a sizeable labor cost shock for cannabis retailers (Saldanha, 2022; Phan, 2022). Further preventative measures include installing panic buttons, two-way door systems, and investing in employee de-escalation training.¹²

3 Data

3.1 Price data

To monitor developments in the cannabis market, legalization came with stringent data reporting and sharing requirements for cannabis businesses. Retailers must record all sales

¹¹Cannabis stores have access to other financial services such as commercial bank accounts (Washington State Department of Financial Institutions, 2022).

¹²Stores are required by law to have comprehensive round-the-clock video surveillance.

and regularly upload data feeds to the LCB. Compliance with data reporting is strictly enforced by the LCB. When a business is issued a violation, it can receive a fine, a temporary license suspension, or both. In cases of repeated violations, a license can be revoked by the LCB board. Given such strict enforcement, violations are uncommon. In 2022 for example, the LCB issued 63 violations among over 1,100 active licensees (Washington State Liquor and Cannabis Board, 2024).

The data, which is reported weekly, contains detailed information on the price and quantity of each product sold by a producer-processor to a retailer, and the subsequent price and quantity of that very same product sold at the retail level. We refer to the price charged by a producer-processor as the *wholesale* price, and the price charged by retailers as the *retail* price. This reflects that producer-processors resemble wholesalers when viewed from the perspective of retailers. The data captures granular product differentiation. For example, in our data, a 1.0-gram package and a 2.0-gram package of the same usable marijuana strain produced by the same producer-processor are treated as different products in the data.

The LCB switched providers for its traceability system in October 2017 and again in December 2021, creating two structural breaks in the price data. Our sample period lies between these breaks and spans March 2018 through December 2021. We obtained the data from Top Shelf Data, a data analytic firm that ingests the raw tracking data from the LCB and matches it with additional product information. The data covers the universe of sales from all 511 active retail establishments during the sample period. Per month, retailers sell approximately 14,500 units and earn about \$280,000 in tax-inclusive sales revenue (see Table 1). Over the entire sample period, the data contain \$5.07 billion in tax-inclusive retail sales. All retail prices and revenues reported in this paper are tax-inclusive.

To estimate pass-through semi-elasticities, we follow previous studies (e.g., Renkin et al., 2022; Leung, 2021; Harasztosi and Lindner, 2019) and use as our dependent variable the natural logarithm of the monthly store-level price index. The log price index measures the price inflation for store j in month t , and is denoted as $\pi_{j,t}$:

$$\pi_{j,t} = \ln I_{j,t}, \text{ with } I_{j,t} = \prod_c I_{c,j,t}^{\omega_{c,j,y(t)}} \quad (1)$$

$I_{j,t}$ is an establishment-level Young price index that aggregates price changes across product subcategories c , where each subcategory is a unique category-unit weight combination. The index weight $\omega_{c,j,y(t)}$ is the revenue share of subcategory c in establishment j during the calendar year of month t .¹³ The dependent variable is equivalent to the first difference of the log of the weighted store price level between month t and $t - 1$.

Store-level indices are common in the literature on retail price movements and carry several advantages over more disaggregated product-level price data. First, retail crime

¹³Price indexes are often constructed using lagged quantity weights. Since product turnover is high at retail stores, lagged weights would limit the number of products used in constructing the price indexes (Renkin et al., 2022). Thus, contemporaneous weights are used.

occurs at the store level, making the store a natural unit of analysis. Second, a store-level index allows the researcher to weigh products by their importance for each store. Finally, entry and exit occur at a much higher frequency for products compared to stores, and a product-level time series would contain frequent gaps. Since the vast majority of cannabis businesses have succeeded at staying in business, the store-level panel is much more balanced. We describe the store-level price index in more detail in Appendix B.

Besides prices, we are also interested in the effect of retail crime on other store-level outcomes. First, we consider the possibility that consumers substitute out of victimized and into nearby rival stores following crime incidents (e.g. due to increased prices or fear of physical harm). To investigate this, we construct a store-level quantity index measuring the monthly percent change in quantity sold. The quantity index is constructed similarly as the price index with index weights based on annual revenue shares (see Appendix B).

It is also possible that stores adjust their wholesale expenditures to offset the costs of crime. Wholesale costs are particularly important since the Cost of Goods Sold (COGS) accounts for 80% of cannabis retailers' variable costs (see Appendix A.1), similar to other retail sectors (Renkin et al., 2022). An advantage in our setting is that we directly observe wholesale prices and quantities at the store-product-month level. We construct a wholesale cost index that captures the month over month percent change in wholesale prices paid by a store. The index weights are based on retailers' annual wholesale expenditures (rather than annual revenue as with the price index).¹⁴

Table 2 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 0.002 to 0.029. 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 can be found in Appendix Figure A.5.

Table 2: Dependent variable descriptive statistics

	Mean	SD	Median
Log price index	-0.001	0.029	0
Log quantity index	-0.008	0.21	-0.012
Log wholesale cost index	1.18e-05	0.002	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.

¹⁴In a robustness check, we construct monthly (rather than annual) expenditure weights which capture potential wholesale substitution patterns on the part of retailers. Results are similar to our main specification that uses annual weights. See Section 6 for details.

3.2 Retail crime data

Similar to other industries, crime at retail cannabis stores is not formally tracked or aggregated by law enforcement agencies. However, on behalf of the market at large, Uncle Ike’s, a retailer and member of Washington’s Craft Cannabis Coalition of over 70 small businesses, maintains a public database of retail crime incidents reported by businesses, law enforcement, and the news media. While the database offers a reasonably accurate list of crimes reported to local police departments, it is not an official census of retail crime at cannabis stores. We view this as unproblematic as each incident in our sample period is cross-referenced with either a police report, police case number, or a news article that references a police investigation.¹⁵ This makes it highly unlikely that treatment is falsely assigned to a store that is, in fact, untreated. Conversely, it is possible that not all crimes are reported to the police or news media and, hence, do not appear in the database. In that case, treated stores would potentially be misclassified as controls, leading to attenuation bias and conservative treatment effect estimates. However, since the control group is much larger than the treatment group in our setting, misclassified controls will be dominated by correctly classified controls. As a result, we expect attenuation bias from misclassified controls to be minimal.

Figure 2 shows the locations of stores with a reported crime during our sample period. There were 74 reported crimes at 57 different stores (46 stores were victimized only once and 11 stores were victimized more than once).¹⁶ Of these crimes, 62 were armed robberies and 12 were burglaries. Approximately two-thirds of these crimes occurred at stores in the Seattle area and the neighboring cities of Tacoma and Bellevue, which roughly corresponds to the share of stores in those cities compared to the rest of the state. In general, crime incidence at cannabis stores reflects the geographic distribution of stores.

There is some evidence of seasonality in crime incidents at cannabis retailers, with more crimes occurring in the fourth quarter compared to other quarters (see Appendix Figure A.4). Over the entire sample period, crime incidents at cannabis retailers increased, which roughly tracks the increase in reported crime in Washington state more generally (see Figure 1).


4 Empirical strategy

We are interested in the effect of retail crime on store-level prices. To identify a causal effect, we exploit the quasi-experimental variation in the time and location of retail crimes using a (stacked) difference-in differences (DiD) estimator. The following subsections describe our empirical strategy in more detail.

¹⁵As an example, the database’s tally aligns perfectly with the Bellevue Police Department’s data on armed robberies at cannabis retailers in the city of Bellevue (Saldanha, 2022).


¹⁶One additional store was victimized during our sample period, but it has no clean control stores so we omit the store from our estimation sample. We describe our criteria for defining clean control stores in Section 4.

Figure 2: Retail cannabis stores in Washington state



(A) Retail stores

Crimes per store
■ 1 ■ >1



(B) Stores with reported crimes

Notes: Panel A illustrates the geographical distribution of cannabis retail stores in our sample. Panel B shows the geographical distribution of victimized stores, with stores reporting one crime incident marked by blue dots and those with at least two reported crime incidents marked by red dots.

4.1 Treatment and control groups

Two treatment groups: victimized and rival stores

We define treatment timing as the store-specific month of a robbery or burglary. We distinguish between two treatment groups based on their exposure to retail crime: i) victimized stores, i.e., all stores that directly experience a robbery or burglary, and ii) rival stores, i.e., non-victimized stores located within a 5-mile radius of a victimized store, excluding any victimized store. The first treatment group allows us to analyze the direct impact of crime on victimized stores' pricing strategies, while the second measures the treatment effect at nearby stores. Such effects at nearby stores could result from demand substitution, where customers shift their demand from victimized stores, e.g. out of fear for their own safety or temporary crime-induced store closures, or from strategic pricing responses. Alternatively, they may stem from cost shocks at rival stores. For instance, upon learning about a crime at a victimized store, nearby retailers might adjust their own expectations of being targeted, leading them to invest in precautionary security measures or additional insurance. In assigning treatment to a larger group of nearby units, our strategy resembles those from the spatial treatment literature (see e.g. Keiser and Shapiro, 2019; Lipscomb and Mobarak, 2017).

We choose the 5-mile radius for two reasons. First, the 5-mile radius is sufficiently small so as to reasonably capture local market characteristics that may influence how rival stores react to a nearby store being victimized. Second, victimized stores have an average of twenty competing stores within this radius, which ensures that a consumer at the average victimized store can choose between a sufficient number of alternative stores. To the extent that crime induces changes in consumer behavior, we view it as plausible that these nearby stores may be affected. Results are robust to adjusting the boundary or using a rank-based rival definition (see Appendix E.3).

Control group: unaffected local markets

We compare both treatment groups to a common control group: stores in unaffected local markets. In defining unaffected local markets, 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.

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.

A major advantage in our setting is that—in addition to retail prices and quantities—we also observe prices and quantities for the universe of vertical transactions between producer-

processors and retailers. This allows us to measure retailers’ spatial sensitivity to competitors’ costs and identify the distance at which retailers no longer compete in prices. We estimate the pass-through of the average wholesale costs of competing stores within a certain distance from store j on store j ’s retail prices. We find that sensitivity to competitors’ unit costs dissipates by the 30-mile mark under a variety of specifications (see Appendix D for more details). This suggests that stores located more than 30 miles from a victimized store will not have a strategic price response to the victimized store’s crime pass-through, and hence serve as valid controls from a strategic pricing standpoint.

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 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.

While potential spillovers require omitting nearby stores from the control group, it is important to ensure that the control group remains comparable to the treated stores to the greatest possible extent. We therefore limit the distance at which a store can serve as a control to 60 miles, ensuring that the stores are situated within the same region of the state. Regional comparability is important because Washington state exhibits significant economic and demographic differences across regions, and approximately two-thirds of retailers’ wholesale purchases occur within the same region (see Appendix Table A.2). Unaffected local markets thus consist of stores located in the concentric “donut” between 30 and 60 miles from a victimized store. We conduct several robustness checks to test the sensitivity of our main results to adjusting the boundaries for unaffected local markets (see Appendix E.2).

Panel A in Table 3 compares pre-treatment characteristics of treated and untreated stores in our study. Columns 1 and 2 show that the average unit price at victimized and rival stores is similar prior to treatment. Victimized stores sell more products per month, have higher monthly revenue, and have more product variety compared to rival stores, though the differences are small and never statistically significant. The similarities between victimized and rival stores prior to crimes support our assertion that the location of crime incidents is conditionally quasi-random. Column 3 summarizes the never-treated group and is based on the entire sample period. Note that never-treated stores may or may not serve as controls

depending on their location. The mean price at never-treated stores is similar to victimized and rival stores, but the quantity sold, monthly revenue, and product variety are slightly lower than at victimized and rival stores, though the differences are not statistically significant.

Table 3: Estimation sample summary statistics

(A) Pre-treatment characteristics			
	(1) Victimized	(2) Rivals	(3) Never-treated
Unit price (in dollars)	26.24 (4.33)	25.47 (4.49)	26.00 (4.50)
Units sold per month	13,260 (12,409)	12,304 (11,931)	10,908 (11,617)
Monthly revenue (in dollars)	248,516 (234,911)	215,201 (217,957)	210,416 (235,890)
Unique products per month	505 (388)	425 (346)	393 (324)
(B) Treatment group sizes			
	(1) Victimized	(2) Rivals	(3) Never-treated
Treated stores	57	242	
Control stores	349	349	152
Total stores	372	410	
Total store-months	30,761	32,249	

Notes: Panel A summarizes store-level variables prior to treatment, with statistics for the never-treated group based on the entire sample period. Standard deviations are in parentheses. The reported variables include mean unit price, mean quantity sold per month, mean revenue per month, and mean number of distinct products sold per month. Columns 1-2 display summary statistics for the two treatment groups: victimized stores and rival stores. Column 3 shows the subset of control stores that are neither victimized nor rival stores. Panel B shows the number of stores in each treatment group.

4.2 Main specification

Our main specification builds on the stacked DiD estimator. Stacked DiD has been applied in a variety of settings, including investigating the effect of local cost shocks on national retail chains (Butters et al., 2022), the effect of application costs on the targeting of disability programs (Deshpande and Li, 2019), and the effects of minimum wage hikes on low wage employment (Cengiz et al., 2019). Stacked DiD overcomes issues with canonical DiD which can yield biased estimates under staggered treatment adoption and heterogeneous treatment effects (see, e.g., Baker et al. (2022); Goodman-Bacon (2021); Sun and Abraham (2021); Callaway and Sant’Anna (2021)). Similar to other DiD estimators, the stacked DiD estimator identifies the causal treatment effect under the assumption of parallel trends and no anticipation. We discuss these identifying assumptions and other threats to identification in more detail in Section 6.

The idea behind stacked DiD is to create a separate dataset (i.e. sub-experiment) for each crime incident comparing the treated group for that incident to a set of “clean” control stores. An identifying variable is generated for each dataset and the datasets are concatenated to create a single “stacked” dataset. The model is estimated on the stacked data with sub-experiment-specific unit- and time-fixed effects. For each sub-experiment, stacked DiD identifies a group average treatment effect for the treated (ATT), similar to the cohort-specific ATT in Sun and Abraham (2021) and the group-time ATT in Callaway et al. (2024); see Section 6 for further details.

Most studies using stacked DiD define clean controls solely based on treatment timing as this accounts for biases from heterogeneous treatment effects under staggered adoption.¹⁷ We add a geographic layer to the clean control criteria to remove additional contamination from treatment effect spillovers. Thus, to be considered a clean control for a given sub-experiment, a store must be (i) located between 30-60 miles of the corresponding crime incident, and ii) a not-yet-treated store outside of its own event window or a never-treated store. A store can qualify as a clean control for multiple sub-experiments with overlapping event windows. Since the store is included in each of the sub-experiment-specific datasets, the stacked dataset will contain duplicate observations of that store for some calendar months—a common feature of the stacked DiD estimator. Further details about our inclusion criteria for clean controls are discussed in Appendix E.1.

The flexibility of the stacked DiD estimator in extending rules for clean controls to geographic criteria is crucial in our setting. In particular, stacked DiD reduces biases from spillovers to untreated stores, maintains geographically comparable treatment and control stores, and ensures a large pool of clean control stores. This feature is the primary reason we choose the stacked DiD estimator over related estimators (Callaway and Sant’Anna, 2021;

¹⁷An exception is Butters et al. (2022), who investigate spillover effects of state excise taxes on chain stores in nearby states. Their approach is similar to ours in that they exclude chain stores whose parent company is affected by a tax hike in another state within the event window.

Borusyak et al., 2024) which typically impose restrictions on the control group composition for the entire sample and thereby significantly reduce the pool of valid control observations.

Panel B of Table 3 summarizes the estimation samples for our victimized and rival specifications. Both specifications have the same control group comprising 349 stores, 152 of which are never-treated. The rival specification (Column 2) has more treated stores than the victimized specification (Column 1), and hence more total stores and more store-months.

We specify a distributed lag model with leads and lags before and after treatment. We define treatment relative to the store’s first crime incident. In our main specification, we drop victimized stores with multiple treatments from the sample prior to the second treatment, isolating the effect of the first crime on store prices (we allow for multiple treatments in a robustness check in Section 6). Following the same logic, we assign treatment for rival stores according to a store’s first nearby crime incident, so that each rival store is assigned to only one treatment cohort. However, we keep a rival store in the sample in subsequent periods even if additional nearby crimes occur. We do this for two reasons: First, several rival stores have multiple nearby crime incidents during their event window. Therefore, dropping these stores prior to the second nearby crime incident would result in an increasingly unbalanced panel at higher lags. Second, to the extent that multiple nearby crime incidents may have a compounding effect on rival stores’ prices, it is important to include these incidents in our analysis. We drop rivals from the sample prior to their second nearby crime incident in a robustness check in Section 6).

Since the store-level price indexes and the treatment adoption indicators are in first-differences, we estimate the following model in first-differences:

$$\pi_{j,t,d} = \sum_{l=-k+1}^k \beta_l \Delta T_{j,d,t-l} + \gamma_{t,d} + \epsilon_{j,d,t} \quad (2)$$

Equation 2 relates the monthly inflation rate at store j , $\pi_{j,t,d}$, to leads and lags of the treatment adoption indicator, $\Delta T_{j,d,t-l}$. $\Delta T_{j,d,t-l}$ is equal to one if store j in sub-experiment d is treated l months after period t (or l months before when l is negative), and equal to zero otherwise. We control for sub-experiment-by-time fixed effects, $\gamma_{t,d}$. As the model is in first differences, store FE are swept out. We cluster standard errors by store to allow for autocorrelation in unobservables within stores, similar to Butters et al. (2022); Deshpande and Li (2019).¹⁸

The parameter β_l measures the change in establishment j ’s inflation resulting from a robbery or burglary l months after a crime (or l months before when l is negative) compared to the control group.¹⁹ While inflation is the dependent variable, we follow previous studies

¹⁸Wing et al. (2024) show that clustering at the level of the treatment (i.e. stores) leads to valid inference with a stacked DiD estimator.

¹⁹Since distributed lag coefficients measure incremental treatment effects, one fewer lead has to be estimated compared to an event study specification. Thus, an 13-month event window requires estimating 12 distributed lag coefficients (Renkin et al., 2022; Schmidheiny and Sieglöcher, 2023).

and present the estimates as the effect of treatment on the price level (see e.g. Renkin et al. (2022); Leung (2021)). We thus normalize the effect on the price level to zero in the baseline period one month before treatment and report the cumulative treatment effect as the sum of β_l at various lags: $E_L = \sum_{l=0}^L \beta_l$. The pre-treatment coefficients are reported in a similar manner with $P_{-L} = -\sum_{l=-1}^{-L+1} \beta_l$. This linear transformation is commonly applied in studies with a dependent variable that is in first-differences (Renkin et al., 2022; Leung, 2021).²⁰ The cumulative distributed lag coefficients are numerically equivalent to the parameter estimates from an event study design under constant treatment effects outside the effect window (Schmidheiny and Siegloch, 2023).

An important consideration is the number of leads and lags to include in equation 2. One limitation is that the establishment panel is not balanced, meaning that changes in the underlying sample may affect estimates when l is large. Therefore, in our baseline estimation, we set $k = 6$. This implicitly assumes constant treatment effects outside of the 13-month event window, i.e. $\beta_l = 0 \forall l > |k|$. We show estimation results using a longer event window in Appendix G.2.

Our main analysis investigates the effects of retail crime on prices at retail cannabis stores. However, we are also interested in the effects of retail crime on other store-level outcomes, as this sheds light on the underlying factors driving retail crime pass-through. For this purpose, we estimate equation 2 using the store-level quantity indexes and wholesale cost indexes as dependent variables.

5 Results


5.1 Main results

Figure 3(A) shows the estimated price level effects for victimized and rival stores according to our main specification. We report cumulative effects, i.e. the effect on the price level (in %) relative to the baseline period in $t - 1$. For both groups, the figure shows no significant pre-treatment effects. In the month of a retail crime incident, however, prices at victimized stores increase and continue to rise for three months. Four months after the crime, prices at victimized stores are 1.8% higher compared to the month before the crime and stay at this higher price level for the following months. In contrast, rival stores show no price increase in the month that a retail crime hits a nearby store. Nevertheless, two months after crime incidents at victimized stores, a very similar price increase can be observed at rival stores. Four months after the crime, rival store prices are 1.5% higher than the month before the crime, and the effect is not statistically significantly different from that at victimized stores. Overall, our estimates reveal that retailers in the Washington state cannabis market increase


²⁰The linear transformation transfers the statistical properties (consistency and asymptotic normality) of $\hat{\beta}$ to the cumulative \hat{E}_L and \hat{P}_{-L} . Standard errors of \hat{E}_L and \hat{P}_{-L} are calculated from the variances and covariances of the vector $\hat{\beta}$ by the standard formula for linear combinations (Schmidheiny and Siegloch, 2023).

their prices following retail crime incidents.


Figure 3: Effect of crime incidents on store prices, quantities sold and wholesale costs



(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. Coefficients are interpretable as percent 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 log price index (Panel A), store-level log quantity index (Panel B), and store-level log wholesale cost index (Panel C).

Figure 3(B) presents estimates from equation 2 when using the quantity index as the dependent variable. Post-incident estimates indicate a 3% (but insignificant) decrease in quantity sold for victimized stores immediately after a crime incident. However, the quantity sold returns to pre-incident levels within a month. This finding speaks against demand substitution out of victimized and into rival stores due to, e.g., stigmatization or personal safety concerns on the part of consumers. Moreover, our findings suggest that demand in our setting is relatively inelastic in the short run. Nevertheless, when we expand the event window, we find a large and statistically significant negative effect on quantity sold after nine months at victimized stores (see Appendix G.2). This delayed effect is consistent with high consumer search costs and persistently higher prices at affected stores (see Appendix Figure

A.14). Pre-incident estimates for victimized stores are insignificant at the 10% level. For rival stores, we observe a slightly negative pre-trend with a statistically significant coefficient at $t-2$. As shown in Appendix G.3, winsorizing the quantity index, which exhibits considerably more volatility than the price index, mitigates this pre-trend. This correction suggests that the observed pre-trend is driven by a small number of outliers in the quantity index.

Figure 3(C) shows estimates with the wholesale cost index as the dependent variable. Estimated monthly treatment effects on wholesale costs are statistically significant but economically small, at approximately 0.1%. These negligible effect sizes imply that retailers do not meaningfully adjust their wholesale expenditures in response to retail crime incidents, and also indicate no major reaction from producer-processors to crime incidents at their client stores. Additionally, these results provide reassurance that changes in COGS are not a confounder driving the observed retail price effects in Figure 3(A). Pre-incident effects show no discernible divergence in the evolution of wholesale costs between the treatment and control groups.

5.2 Heterogeneity analysis

In this subsection, we conduct two main heterogeneity analyses to gain more insight into the role of market structure in crime pass-through.

Chain vs Independent stores

We first examine whether crime pass-through differs between chain and independent stores. One concern is that chain size may correlate with profitability and is therefore endogenous. This is less of a concern in our setting for two reasons. First, when the market was initially created, the LCB randomly assigned licenses to applicants, reducing self-selection bias for store types Hollenbeck and Giroldo (2022). Second, the number of licenses is capped by the LCB, meaning that profitable firms cannot easily increase their chain size.

We define chains as stores belonging to firms running three or more stores and consider all other stores as independent. Of the 57 victimized stores, 37 are independent and 20 belong to chains. We sort sub-experiments into two groups depending on whether the corresponding victimized store is a chain store or independent. For each sub-experiment, we preserve the original control group (results are similar if we also subset control groups by chain size). We then estimate our main specification separately for each group. For the rival specification, we sort rival stores into two subgroups depending on whether a rival store is independent or belongs to a chain (of the 242 rival stores, 154 are independent and 88 belong to chains). We then estimate our rival specification separately for each subgroup (using the original control groups).

Table 4 reports the results of the subsample analyses. Columns 1 and 2 show that among victimized stores, the pass-through of crime is large at independent stores but smaller and

not statistically significantly different from zero at chain stores. Columns 5 and 6 show that the same pattern holds among rival stores. When defining independent as single-store firms and chains as firms with two or more stores (see Appendix Table A.5), we observe a similar difference for victimized stores but not for rival stores. Overall, our findings align with studies showing little within-chain price adjustments to local conditions (Hitsch et al., 2021; DellaVigna and Gentzkow, 2019) and suggest that owning multiple stores may serve as insurance against negative cost shocks, possibly due to better financing conditions or high returns to scale (Hollenbeck and Giroldo, 2022). Alternatively, or as an additional mechanism, the results suggest that chains may already have more effective crime prevention strategies or insurance plans in place that require less adjustment, supporting the notion of firm learning.

Market concentration

Next, we investigate whether effects vary by market concentration. Concentration is endogenous if profitability affects market entry. However, since the LCB caps the number of cannabis store licenses and distributes them according to population density, profitability does not directly affect concentration in our setting. We consider each victimized store as the focal point of its own market comprising the set of stores (including the focal store) within a 5-mile radius. We calculate the Herfindahl–Hirschman Index (HHI) for each market and sort sub-experiments into two groups depending on whether the victimized store’s market HHI is above or below the median HHI for all victimized store markets. A high HHI implies high concentration and thus low market competition. We provide descriptive statistics on the HHI in Appendix Table A.4.

We estimate our main specification on each subsample to test whether crime pass-through differs between stores above and below the median HHI. Columns 3 and 4 of Table 4 show that among victimized stores, crime causes prices to increase 3.1% in low-concentration markets while in high-concentration markets the effect is not significantly different from zero. Moreover, the difference in effect sizes is statistically significant. This pattern carries over to rival stores, though the difference in effects sizes between low and high market concentration is less pronounced. Prices in low-concentration markets increase 2.2% while in high-concentration markets they increase 1.1%.²¹

These results are in line with standard imperfect competition models (e.g. the canonical Cournot model), which predict that when competition increases, prices become more sensitive to marginal costs and pass-through rates rise. However, if marginal cost pass-through rates exceed unity or the shape of the cost function changes, predictions can change (Weyl and Fabinger, 2013; Miller et al., 2017; Ritz, 2024). Empirical findings linking competition and

²¹In Appendix C, we investigate whether differences between rural and urban markets explain the heterogeneous pass-through rates. Our analysis indicates that once we account for variations in market concentration between rural and urban areas, the disparity in effect sizes between markets with low and high concentration largely diminishes (Appendix Table A.5).

cost pass-through rates are similarly mixed. Some studies report similar findings to ours, showing that pass-through rates increase with competition (Cabral et al., 2018; Genakos and Pagliero, 2022), while other studies find the opposite (Doyle Jr and Samphantharak, 2008; Miller et al., 2017).

Table 4: Heterogeneity of retail crime pass-through

	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
E_0	0.010 (0.0079)	0.0018 (0.0026)	0.011 (0.010)	0.0041 (0.0030)	0.00071 (0.0026)	0.0022 (0.0030)	0.000077 (0.0025)	0.0029 (0.0024)
E_2	0.021** (0.0096)	0.0037 (0.0044)	0.022* (0.012)	0.0086 (0.0056)	0.0077* (0.0040)	-0.00088 (0.0045)	0.0048 (0.0044)	0.0046 (0.0033)
E_4	0.026** (0.012)	0.0035 (0.0051)	0.031** (0.014)	0.0056 (0.0077)	0.021*** (0.0061)	0.0052 (0.0068)	0.017*** (0.0065)	0.011* (0.0062)
E_6	0.026** (0.013)	0.016 (0.012)	0.034** (0.014)	0.012 (0.012)	0.019*** (0.0071)	-0.0018 (0.0078)	0.014* (0.0074)	0.0060 (0.0080)
\sum Pre-event	0.00035 (0.0079)	-0.0015 (0.0060)	0.0040 (0.0080)	-0.0046 (0.0077)	0.0018 (0.0067)	0.0041 (0.011)	-0.0024 (0.0069)	0.0094 (0.0070)
N	21,037	9,724	14,826	15,935	29,956	28,837	17,974	13,050

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 percent 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 (i.e., owned by firms running one or two stores only), while columns 2 and 6 show effects for stores owned by firms running at least three stores. For the other columns, the Herfindahl-Hirschman Index is calculated for each victimized store’s local market market, defined as all stores within a 5-mile radius of the victimized store. The sample is split according to the median market concentration across all victimized markets. Columns 3 and 7 show effects for treated stores in 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6 Threats to identification, robustness checks and alternative specifications

In this section, we address potential endogeneity concerns in our setting. We present various alternative specifications, robustness checks and placebo tests to corroborate the validity of our results.

Endogenous treatment timing and parallel trends

The key identifying assumption of our empirical strategy is the parallel trends assumption. This implies that store-level prices in the control and treatment groups would have followed a common trend in the absence of retail crime incidents.

In our framework, we account for all time-invariant factors that could influence treatment, such as store location and average revenues, through the inclusion of store-fixed effects. Similarly, we account for all time-varying factors that equally apply to all stores, such as the seasonality of robberies or COVID-19 effects, through month-year fixed effects. However, the parallel trends assumption can be violated if the timing and location of retail crimes are correlated with changes in our outcome variables. An example is if stores that strongly increase or decrease revenues are more likely to be robbed or burglarized. This is less of a concern with rival stores' treatment timing. However, it is also possible that changes in policing in certain areas are correlated with revenues in those areas, which could also threaten the causality for our rival store regression.

The first clear indication of the validity of the parallel trends assumption in our setting is the lack of significant pre-trends across all main specifications. Figure 3 shows no pre-treatment differences for prices, quantity sold and wholesale cost between victimized, rival, and control stores. Furthermore, pre-trends remain insignificant for alternative specifications, for instance, if we expand the event window (Appendix Figure A.14) or when using alternative estimators (Appendix F). Additionally, the lack of observable pre-trends supports the validity of the second identifying assumption regarding no anticipation. If store owners anticipate robberies, we would expect to see changes in the outcome variables before the actual events.

To further corroborate our findings, we conduct placebo tests in which we run our main regression analysis with a placebo treatment date that occurs 12 months before the actual treatment date. This approach tests for the presence of non-parallel trends or seasonal variations not fully addressed by our fixed effects. The placebo test results, detailed in Appendix Figure A.13, reveal no discernible patterns, reinforcing our confidence in the validity of the parallel trends assumption and our empirical model.

Alternative estimators

The stacked DiD estimator offers advantages over related estimators due to its flexibility for applying rules for clean controls based on geographic criteria (for more details, see Appendix F). Nevertheless, other estimators may have advantages in terms of efficiency and comparability to other studies. To assess whether our findings are sensitive to the choice of estimator, we estimate dynamic treatment effects using three alternative estimators: i) the canonical two-way fixed effect DiD estimator ii) the imputation estimator developed by Borusyak et al. (2024) (BSJ); and iii) the estimator proposed by Callaway and Sant'Anna (2021) (CS).²²

²²The canonical TWFE estimator is prone to biases under staggered treatment adoption. The BSJ and CS estimators address these biases in a different way than the Stacked DiD estimator.

Following our main approach, we first estimate a distributed lag model (Equation 2) for the designated event window and then calculate cumulative treatment effects using the last pre-treatment period as our baseline. We discuss the details of the alternative estimators in Appendix F.

The dataset used for alternative estimators includes all stores, which means the control group includes stores that are geographically close to victimized stores. Positive price effects at these nearby stores, as suggested by our main results, would thus imply that treatment effects are attenuated in these models. Restricting the sample to stores more than 30 miles away from any treated store reduces the number of control stores from 349 to just 78 stores, reflecting that most stores are within 30 miles of at least one retail crime incident at some point during the sample period. This highlights the benefits of using the stacked DiD framework in our setting.

Appendix Figure A.12 displays the cumulative treatment effects for the alternative estimators. As expected, the effects are slightly smaller than in our main specification—about 1.5% for victimized stores and 1% for rival stores, except for the CS specification, which yields higher effects for victimized stores. However, in the majority of these alternative specifications, treatment effects are statistically significant and not different from our main specification, demonstrating that our findings are not driven by the choice of our estimator.

Alternative definitions of clean control and rival stores

When defining clean control stores, we balance comparability with treated stores against potential biases from treatment spillovers, as discussed in Section 4. To investigate whether our estimated effects are sensitive to our definition of clean control stores, we estimate our model using alternative definitions unaffected local markets and rival stores. Moreover, we investigate the sensitivity of our results to using more restrictive inclusion criteria.

First, we expand the area of unaffected local markets along two dimensions: i) including all stores between 10 and 60 miles of victimized stores; ii) including all stores that are located more than 30 miles away from victimized stores. Results presented in Appendix E.2 show that treatment effect estimates are very similar and statistically significant when including additional control stores closer to victimized stores. Extending the control group to include more distant stores results in slightly smaller treatment effects—around 1.3% for victimized stores and 0.8% for rival stores four months post-crime incident—yet these effects still retain statistical significance for both groups.

Next, we inspect the sensitivity of our results to expanding the inner ring that defines rival stores. We report results in Appendix Figure A.8(C). As expected, the treatment effect size decreases when considering stores further away from victimized stores as rivals, indicating that spillovers decrease with distance from the crime incident. This decline supports our choice of control store as it suggests that treatment spillovers affect stores located further away less strongly. Nonetheless, the treatment effect remains sizeable at around 1% four

months after an incident for stores within 10 miles and only becomes statistically insignificant for those further away. These results imply that spillover effects may extend to a larger set of stores than initially anticipated and considered in our welfare analysis.

Our definition of unaffected local markets and rival stores is constant across locations, but market characteristics may vary depending on whether a store is located in an urban or rural area. For instance, urban stores typically have more nearby rival and control stores, and urban customers may face higher transportation costs when traveling longer distances. To address potential biases related to these differences, we run two alternative specifications, detailed in Appendix E.3. First, we redefine unaffected local markets by using distance ranks from the victimized store, defining the 150th to 250th closest stores as clean controls and the 20 closest stores as rivals. This approach balances the number of control and rival stores between urban and rural sub-experiments. Second, we tighten the boundaries for defining clean control stores for victimized stores in urban areas, using a 10-30 mile range. The results from these specifications align with our main findings, providing further support for the robustness of our results.

Our main specification excludes victimized and rival stores from serving as clean control stores in another sub-experiment with an overlapping event window. However, untreated stores within 30 miles of a soon-to-be-victimized or recently victimized store can still serve as clean control stores if all inclusion criteria are met. While this permits a large pool of control stores, it introduces potential biases if treatment spillovers affect these untreated stores. In Appendix E.4, we discuss this issue in detail and show that our main results remain robust when applying more stringent clean control criteria to account for potential spillover effects.

Finally, some rival stores experience multiple crime incidents within a 5-mile radius in a single month, which can create variations in treatment intensity and lead to a violation of the SUTVA in our rival store specifications. In Appendix E.5, we show that our main results are robust to the exclusion of rival stores that encounter more than one nearby crime incident within a given month.

Additional Robustness checks

We conduct a number of additional robustness checks to rule out other factors potentially driving our findings. We present the results of these robustness checks in Table 5. In both panels, Column 1 reproduces our baseline specification from Figure 3(A). Column 2 shows that our results are similar when we include control variables such as county population, the local house price index (at the three-digit zip code level), and average county wage. These control variables absorb variation in prices stemming from local business cycles, population growth, or fluctuations in the housing market (Stroebel and Vavra, 2019).

Since our price indexes are store-level aggregations of diverse sets of products and product categories, it is important to check whether results hold for a narrower set of homogeneous products. Therefore, in Column 3 we restrict the sample to products belonging to the usable

marijuana product category and convert all prices to price per gram.²³ Effect sizes are larger using this price index. In Column 4 we winsorize inflation rates below the bottom 0.5 percentile and above the 99.5 percentile to show that our results are not driven by outliers. We further discuss the winsorized specification, including when the quantity index is used as the dependent variable, in Appendix Section G.3.

Our price indexes are constructed with weights based on calendar year revenue shares (see Section 3). To ensure that our main effects are not an artifact of this weighting scheme, in Column 5 we use indexes constructed with weights based on the fiscal year running from July through June. Estimates are similar to our baseline specification.

Our main analysis only includes a store’s first reported crime incident and hence does not capture the effects of subsequent crimes at stores that are victimized more than once (of the 57 victimized stores in our sample period, 11 stores have multiple crime incidents). In Column 6 Panel A, we include all subsequent crimes with non-overlapping event windows, which allows for individual stores to be treated more than once.²⁴ For victimized stores, treatment effects are slightly higher shortly after crime incidents but become attenuated at later lags when allowing for multiple treatments. This is in line with the idea that store owners invest more following the first event but react less to subsequent events.

In the main rival specification, we retain rival stores in the sample for subsequent periods, even if additional nearby crimes occur. We do this to maintain a balanced panel at higher lags and account for potential compounding effects of nearby crimes on rival stores’ prices. In Column 6 of Panel B, we drop rival stores from the sample before their second nearby crime incident to isolate the effect of the first rival treatment. While the effects remain highly statistically significant, dropping rivals before their second incident slightly attenuates the observed price effects at these stores. This attenuation suggests that nearby crime incidents have a compounding effect on rival stores.

By excluding recently or soon-to-be-treated stores from the control group, stacked DiD identifies a group ATT for each crime incident under the standard DiD assumptions of parallel trends and no anticipation, similar to the cohort-specific ATT in Sun and Abraham (2021) and the group-time ATT in Callaway et al. (2024). When there is a balance in the number of pre- and post-treatment periods across sub-experiments, the stacked DiD regression recovers an aggregate ATT that is a convex combination of underlying causal effects. However, even when the DiD assumptions hold within each sub-experiment, the aggregated parameter estimate can be biased if the stacked dataset is not balanced in the number of pre- and post-periods across sub-experiments (Wing et al., 2024). This is the case if some crime incidents occur near the beginning or end of the period for which data are available so that causal effects are

²³All products in the Usable Marijuana product category contain information on the package weight measured in grams. This enables us to convert prices into prices per gram for this product category.

²⁴If the first and subsequent crimes for a given store have overlapping event windows, stacked DiD can produce biased estimates if treatment effects are heterogeneous over time. This is similar to the bias that can arise from staggered DiD.

identified for a larger number of event-time periods for some crime incidents than others.²⁵ The bias arises because changes in the aggregate parameter may reflect compositional changes rather than treatment effect dynamics.

To account for this, in Column 7 we estimate our stacked DiD regression using the subset of crime incidents for which the entire event window falls within our sample period.²⁶ In both panels of Table 5, Column 7 shows that imposing compositional balance in pre- and post-treatment periods substantially reduces our sample size but does not meaningfully change, and even slightly increases, our estimates. Standard errors tend to be larger due to the reduced sample size, but the effects are still significant at the 5% level four months after the crime.

Stacked DiD implicitly weights treatment and control trends differently across sub-experiments. This can lead to bias in the aggregate parameter estimate if the share of treated and control stores differs across sub-experiments. The bias can arise even if the incident-specific ATTs are unbiased and there is compositional balance in pre- and post-periods across sub-experiments. Wing et al. (2024) show that this bias can be corrected by using sample weights that account for relative treatment and control group shares. Using the balanced stacked dataset (from Column 7), we estimate a weighted least squares regression using sample weights based on relative treatment shares in each sub-experiment. We report results in column 8 of Table 5. Effect sizes are comparable and not statistically significantly different from our baseline estimates.

One other concern is that retail crime may cause victimized stores to go out of business. If crime-induced store failure is correlated with unobserved store characteristics (e.g. profitability), then our treatment effect estimates may be biased. However, as detailed in Appendix G.4, we find no significant difference in failure rates between victimized and non-victimized stores. In fact, 5.3% of victimized stores close within a year after an incident, which is slightly lower than the 5.5% average annual closure rate for non-victimized stores. These numbers suggest that crime-induced store closures are an unlikely source of bias.

7 Policy analysis—Retail crime as a hidden tax

We model retail crime as a marginal cost shock that functions like a hidden tax—an analogy first introduced by Jackson and Tran (2020). The marginal cost shock could arise from multiple sources. On the one hand, the shock may represent direct marginal costs proportional to the quantity sold. Industry reports suggest that retailers respond to crime by hiring

²⁵Since our setting contains several crime incidents near the beginning and end of the sample period, dropping these incidents substantially reduces our treatment sample size. Therefore, in our main specification we allow for compositional imbalance in pre- and post-treatment periods across sub-experiments.

²⁶In addition, we only include treated stores for which scanner data is reported for at least 75 % of the event window. We prefer this cutoff because it allows for occasional idiosyncratic store-level gaps in data reporting in the LCB traceability system while at the same time ensuring stability in the composition of treatment and control groups within sub-experiments.

Table 5: Price effects at victimized and rival stores - robustness checks

(A) Victimized stores								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Baseline	Controls	Price per gram	Wins.	Alt. weights	Mult. treat- ments	Bal- anced	Bal- anced (weighted)
E_0	0.007 (0.005)	0.007 (0.005)	0.009 (0.007)	0.008 (0.005)	0.006 (0.004)	0.010 (0.007)	0.011 (0.008)	0.011 (0.008)
E_2	0.015** (0.007)	0.015** (0.007)	0.024*** (0.008)	0.016** (0.007)	0.012* (0.007)	0.014** (0.007)	0.022** (0.009)	0.021** (0.009)
E_4	0.018** (0.008)	0.018** (0.008)	0.028*** (0.010)	0.018** (0.008)	0.015* (0.009)	0.015 (0.011)	0.026** (0.011)	0.025** (0.011)
E_6	0.022** (0.009)	0.022** (0.009)	0.033*** (0.012)	0.022** (0.009)	0.023** (0.011)	0.017* (0.010)	0.022* (0.012)	0.021* (0.012)
\sum Pre-event	0.000 (0.006)	0.000 (0.006)	-0.004 (0.008)	-0.001 (0.006)	0.002 (0.005)	-0.005 (0.005)	0.002 (0.007)	0.002 (0.007)
N	30,761	30,761	30,651	30,761	30,761	31,578	20,793	20,793
(B) Rival stores								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Baseline	Controls	Price per gram	Wins.	Alt. weights	Single treat- ment	Bal- anced	Bal- anced (weighted)
E_0	0.001 (0.002)	0.001 (0.002)	0.005 (0.004)	0.001 (0.002)	0.002 (0.002)	0.002 (0.002)	0.002 (0.003)	0.001 (0.003)
E_2	0.005 (0.003)	0.004 (0.003)	0.007 (0.005)	0.004 (0.003)	0.003 (0.003)	0.006* (0.003)	0.006 (0.004)	0.005 (0.004)
E_4	0.015*** (0.005)	0.014*** (0.005)	0.019*** (0.006)	0.013*** (0.004)	0.013*** (0.005)	0.014*** (0.005)	0.015** (0.006)	0.013** (0.006)
E_6	0.012** (0.006)	0.011* (0.006)	0.014* (0.008)	0.010** (0.005)	0.009 (0.006)	0.009 (0.006)	0.012* (0.007)	0.010 (0.006)
\sum Pre-event	0.003 (0.006)	0.003 (0.006)	-0.001 (0.009)	0.003 (0.005)	0.009* (0.005)	0.002 (0.006)	-0.001 (0.008)	-0.002 (0.008)
N	31,024	31,024	30,891	31,024	31,024	25,877	17,913	17,913

Notes: Each column shows the cumulative treatment effects on store price levels for different specifications zero, two, four and six months after a crime, along with the sums of pre-treatment coefficients. Coefficients are interpretable as percent increases in outcome levels relative to the month before a crime incident. Panel A uses victimized stores as the treatment group, and Panel B considers rival stores. Column 1 in each panel shows the cumulative treatment effects from the main specification. Column 2 includes additional time-varying control variables: county population, local house price index (at the three-digit zip code level), and average county wage. Column 3 restricts the sample to Usable Marijuana products and converts all prices to price per gram. Column 4 uses store-level price indexes winsorized at the 0.5 and 99.5 percentiles as the dependent variable. Column 5 uses price indexes weighted based on the fiscal year running from July through June as the dependent variable. Column 6 (Panel A) considers subsequent store-level crimes as additional treatment events. Column 6 (Panel B) drops rival stores from the sample the month before their second rival treatment. Column 7 shows effects for events with a completely observable and balanced event window. Column 8 applies the sample weights proposed by Wing et al. (2024) to the balanced stacked data set. Standard errors of the sums are clustered at the store level and shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

additional security personnel, expanding surveillance, or increasing theft insurance coverage, all of which scale with store size, operating hours, and product quantities (U.S. Chamber of Commerce, 2023; Dunham, 2021; National Retail Federation, 2022; U.S. Department of Homeland Security, 2022). Alternatively, the shock could arise from changes in retailers’ subjective perceptions of marginal costs. For example, retailers may use average costs as a proxy for marginal costs, meaning fixed costs such as property loss or one-time security expenditures could indirectly influence subjective marginal costs and, in turn, pricing strategies. In our welfare analysis, we remain agnostic about the exact sources of the marginal cost shocks, as the welfare implications generally hold irrespective of their underlying source.

In Section 7.1, we adopt the symmetric oligopoly model of Weyl and Fabinger (2013)—which nests common imperfect competition models as well as the monopoly and perfect competition cases—to derive the general welfare implications of a hidden crime tax. Section 7.2 examines how retailers adjust their prices in response to changes in marginal cost, utilizing methodologies from the extensive literature on cost pass-through (e.g. Hollenbeck and Uetake, 2021; Muehlegger and Sweeney, 2022; Conlon and Rao, 2020; Miller et al., 2017). Our objective in estimating the marginal cost pass-through rate is twofold. First, assuming that crime-related cost shocks are passed through at a similar rate as other marginal cost shocks, the marginal cost pass-through rate enables us to quantify the hidden crime tax. Second, our analysis of marginal cost pass-through encompasses not just a retailer’s own marginal costs but also the average marginal costs of its competitors, which allows us to examine strategic complementarity in pricing in the cannabis industry. The degree of strategic price complementarity is informative about the mechanisms driving rivals’ price response to crime and has important implications for our welfare analysis. In Section 7.3, we combine the estimated marginal cost and retail crime pass-through rates with the theoretical insights from Weyl and Fabinger (2013) to derive and quantify the welfare effects of the hidden crime tax in our context.

7.1 Welfare implications 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).²⁷ In this case, Weyl and Fabinger (2013) show that under

²⁷This assumption is supported by our main results and our marginal cost pass-through estimates (see Section 7.2). 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.

symmetric imperfect competition the pass-through rate (in dollars) for a small unit tax is:

$$\rho = \frac{dp}{d\tau} = \frac{1}{1 + \frac{\epsilon_D - \theta}{\epsilon_S} + \frac{\theta}{\epsilon_{ms}} + \frac{\theta}{\epsilon_\theta}}, \quad (3)$$

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 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:

$$\frac{dCS}{d\tau} = -\rho q \quad (4)$$

$$\frac{dPS}{d\tau} = -[1 - \rho(1 - \theta)] q \quad (5)$$

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:

$$I = \frac{\rho}{1 - \rho(1 - \theta)} \quad (6)$$

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.

7.2 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). A major advantage of our setting is that, because we observe wholesale unit prices, we can directly measure how changes in marginal cost are passed through to prices.²⁸

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):

$$\Delta p_{i,j,t} = \rho \Delta w_{i,j,t} + \beta \Delta w_{i,r(j),t} + \gamma_t + \Delta \varepsilon_{i,j,t}, \quad (7)$$

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.²⁹ 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 D, 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 6 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

²⁸Wholesale costs are typically estimated from supply-side first order conditions as such granular information on vertical relations is rarely observable by researchers (Dubois et al., 2022).

²⁹Since 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).

(2021), and indicates that cannabis retailers exercise substantial market power.³⁰ According to equation 3, 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 robustness checks in Appendix D, 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 6 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 6 and Appendix D.

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.

7.3 Quantifying the welfare effects of retail crime pass-through in cannabis

To quantify the hidden tax of crime and the corresponding welfare effects in the Washington state cannabis industry, we combine the sufficient statistics approach outlined in Section 7.1, our DiD estimates from Section 5, and our marginal cost pass-through estimates from Section 7.2. Our goal is not to measure the overall welfare effects of retail crime in the Washington state cannabis market, as these include other direct and indirect costs of crime that we do not observe. Instead, we seek to quantify the welfare effects specifically stemming from the pass-through of retail crime onto stores' prices.

As a first step, we quantify the hidden unit tax from crime. In Section 5, we found that stores in affected markets (i.e., victimized and rival stores) raise prices by 1.5% four months after a crime incident. Multiplying this semi-elasticity by the mean unit price of

³⁰Pass-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).

Table 6: Unit cost pass-through rates

	(1) Baseline	(2) Controls	(3) Reg. \times time FE
Own wholesale cost	1.67*** (0.027)	1.67*** (0.027)	1.67*** (0.027)
Rivals' wholesale cost (0-5 miles)	0.016*** (0.006)	0.017*** (0.006)	0.016*** (0.06)
N	5,260,269	5,245,497	5,260,269

Notes: The table reports the estimates of wholesale unit cost pass-through rates from equation 7. 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

\$27.93 in affected markets, we find that crime causes an increase in the average unit price of around \$0.42. We then compare this to the marginal cost pass-through rate obtained in Section 7.2. Assuming that crime-related changes in marginal cost are passed through at a similar rate, we can calculate the implied hidden unit tax of crime by dividing our treatment effect in dollar terms by the pass-through rate.³¹ This yields a unit tax per crime incident of about \$0.25—or approximately 1% at the average unit price. Next, we calculate the fictional annual tax revenue. Depending on the source of the marginal cost shock, this fictional revenue could represent increased expenditures on security or insurance. If the marginal cost shock is subjective, the fictional tax revenue would instead reflect unanticipated future profits for retailers. The fictional tax revenue can be approximated by multiplying the unit tax per crime (\$0.25) by the average annual units sold in affected local markets. The average annual quantity sold in affected markets is 45,520,552 units, implying a fictional tax revenue of around \$11.4 million per year.

As a next step, we calculate the effect on consumer surplus by multiplying the crime-induced increase in unit prices (\$0.42) by the average annual units sold in affected markets. This yields a negative annual effect on consumer surplus of around \$19.1 million. To quantify the effect on producer surplus and the incidence of the hidden tax, we rely on the conduct parameter estimate by Hollenbeck and Uetake (2021), who study the same industry using comparable data. They estimate an average θ of 0.89, which aligns well with our results of substantial market power by retailers. Together with our marginal cost pass-through estimate and the average annual quantity sold, equation 5 implies that a \$1 unit tax reduces producer surplus by around \$37.2 million, or a decrease of \$9.3 million for the estimated hidden crime

³¹There is no scope for factor substitution in cannabis retail, meaning the marginal cost pass-through rate should be independent of the source of the marginal cost change (Ganapati et al., 2020).

tax.³² Accordingly, the hidden crime tax incidence falls about 67% on consumers and 33% on producers. The incidence of the tax on consumers is equal to around 62.5% in the monopoly case ($\theta = 1$) and increases if retailers' market power decreases.

The combined annual welfare effect for market participants is about \$28.4 million. Assuming that the fictional tax revenue stays within the Washington state economy, we find an excess burden from retail crime pass-through of around \$17.0 million per year (\$28.4 million - \$11.4 million). Note that these welfare calculations assume that only stores within a 5-mile radius of an incident are affected. Further analysis presented in Appendix E.2 indicates that while the effect of retail crime on prices decreases for stores located further away, it remains significant. This finding suggests that our welfare analysis is conservative, indicating that the actual welfare impacts on Washington state's economy could be higher than estimated.

8 Discussion and conclusion

In this paper, we present the first causal evidence that retailers pass costs related to organized retail crime through to consumer prices, uncovering an overlooked aspect of retail crimes' social costs. We show this using novel matched administrative datasets from the Washington state cannabis industry which provides an ideal setting to estimate causal relationships and disentangle the channels driving the pass-through of retail crime.

While the cannabis industry is distinct, our findings have significant welfare implications beyond our context. The cannabis industry shares several important characteristics with other retail sectors. These include a comparable variable cost structure (see Appendix A.1) and similar demand elasticities, both of which are important determinants of cost pass-through (Hollenbeck and Uetake, 2021). The lack of a competing black market (Hollenbeck and Uetake, 2021) underscores the generalizability of our results. Further support for the broad applicability of our findings comes from a variety of news reports, a recent survey by the National Retail Federation (2022), and from findings of a study by Jackson and Tran (2020), who show that a decreased likelihood of a felony conviction is connected to increased prices for automobiles and computers. These sources highlight that increased crime-related costs are a widespread concern for retailers in numerous sectors, suggesting that these costs are passed on to consumers in a variety of settings.

We estimate an annual excess burden for the Washington state cannabis industry of around \$17.0 million. In 2020, the state's total cannabis sales amounted to \$1.4 billion, while total U.S. retail sales reached \$5,572 billion (U.S. Census Bureau, 2022). When we scale the excess burden of retail crime pass-through to all U.S. retailers based on the relative shares of sales, our estimates imply an additional annual welfare cost of about \$68 billion nationwide. This naive extrapolation requires careful interpretation since the impact of

³²Note that the reduction in producer surplus entails a decrease in quantity demanded. While our main results reveal no change in quantity sold following a retail crime incident, we find evidence of a delayed reduction in quantity sold at higher lags (see Appendix G).

retail crime on prices likely differs across sectors due to different market structures and crime rates. Nevertheless, even if the pass-through rate of retail crime in other sectors is substantially lower, our analysis underscores the importance of considering pass-through effects when evaluating the social costs of retail crime. Furthermore, with state and local governments allocating approximately \$126 billion to police protection in 2020 (Moore et al., 2022), our findings provide a compelling argument for increasing public expenditure on crime reduction.

Our welfare analysis makes some restrictive assumptions on the implications of retail crime pass-through. For instance, we abstract from the potential negative external effects of cannabis consumption. While the cannabis sales tax aims to internalize these external effects, the identified price hikes may still have beneficial corrective effects for some consumers. Moreover, we neglect the distributional implications of our results. The price effects may be regressive in the sense that they fall disproportionately on low-income and younger individuals, aligning with the demographics of cannabis consumers (see Appendix A.1). A detailed analysis of these additional welfare implications goes beyond the scope of this paper and is left for future research.

References

- Audretsch, D. B. and Feldman, M. P. (2004). Knowledge spillovers and the geography of innovation. In *Handbook of Regional and Urban Economics*, volume 4, pages 2713–2739. Elsevier.
- Axios (2024). Americans’ average daily travel distance, mapped. <https://www.axios.com/2024/03/24/average-commute-distance-us-map>. Retrieved April 24, 2024.
- Baker, A. C., Larcker, D. F., and Wang, C. C. (2022). How much should we trust staggered difference-in-differences estimates? *Journal of Financial Economics*, 144:370–395.
- Bloom, N., Eifert, B., Mahajan, A., McKenzie, D., and Roberts, J. (2013). Does management matter? Evidence from India. *Quarterly Journal of Economics*, 128:1–51.
- Borusyak, K., Jaravel, X., and Spiess, J. (2024). Revisiting event study designs: Robust and efficient estimation. *Review of Economic Studies*, pages 1–33.
- Butters, R. A., Sacks, D. W., and Seo, B. (2022). How do national firms respond to local cost shocks? *American Economic Review*, 112:1737–1772.
- Cabral, M., Geruso, M., and Mahoney, N. (2018). Do larger health insurance subsidies benefit patients or producers? Evidence from medicare advantage. *American Economic Review*, 108(8):2048–2087.
- Callaway, B., Goodman-Bacon, A., and Sant’Anna, P. H. (2024). Difference-in-differences with a continuous treatment. *NBER Working Paper 32117*.
- Callaway, B. and Sant’Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225:200–230.
- 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.
- Chapekis, A. and Shah, S. (2024). Most americans now live in a legal marijuana state – and most have at least one dispensary in their county. Technical report, Pew Research Center.
- 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.
- Cullen, J. B. and Levitt, S. D. (1999). Crime, urban flight, and the consequences for cities. *Review of Economics and Statistics*, 81(2):159–169.
- De Chaisemartin, C. and d’Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9):2964–2996.

- De Mello, J. M. and Zilberman, E. (2008). Does crime affect economic decisions? An empirical investigation of savings in a high-crime environment. *The BE Journal of Economic Analysis & Policy*, 8(1).
- DellaVigna, S. and Gentzkow, M. (2019). Uniform pricing in U.S. retail chains. *The Quarterly Journal of Economics*, 134(4):2011–2084.
- Deshpande, M. and Li, Y. (2019). Who is screened out? Application costs and the targeting of disability programs. *American Economic Journal: Economic Policy*, 11:213–248.
- Doyle Jr, J. J. and Samphantharak, K. (2008). \$2.00 gas! Studying the effects of a gas tax moratorium. *Journal of Public Economics*, 92(3-4):869–884.
- Dubois, P., Griffith, R., and O’Connell, M. (2022). The use of scanner data for economics research. *Annual Review of Economics*, 14(1):723–745.
- Dunham, J. (2021). The impact of organized retail crime and product theft in the United States. Technical report, Retail Industry Leaders Association and Buy Safe America Coalition.
- Dyer, J. (2023). The fruits (and vegetables) of crime: Protection from theft and agricultural development. *Journal of Development Economics*, 163:103109.
- Fe, H. and Sanfelice, V. (2022). How bad is crime for business? Evidence from consumer behavior. *Journal of Urban Economics*, 129:103448.
- Fenzia, A. and Saggio, R. (2024). Organized crime and economic growth: Evidence from municipalities infiltrated by the mafia. *American Economic Review (Forthcoming)*, (w32002).
- Fonrouge, G. (2022). Rising thefts at walmart could lead to price jumps, store closures, CEO says. <https://www.cnbc.com/2022/12/06/walmart-ceo-says-shoplifting-could-lead-to-price-jumps-store-closures.html>. Retrieved December 17, 2023.
- 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.
- Genakos, C. and Pagliero, M. (2022). Competition and pass-through: Evidence from isolated markets. *American Economic Journal: Applied Economics*, 14(4):35–57.
- Gibbons, S. (2004). The costs of urban property crime. *The Economic Journal*, 114(499):F441–F463.
- Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225:254–277.

- Greenbaum, R. T. and Tita, G. E. (2004). The impact of violence surges on neighbourhood business activity. *Urban Studies*, 41(13):2495–2514.
- Hamermesh, D. S. (1999). Crime and the timing of work. *Journal of Urban Economics*, 45(2):311–330.
- Harasztosi, P. and Lindner, A. (2019). Who pays for the minimum wage? *American Economic Review*, 109:2693–2727.
- Hipp, J. R., Williams, S. A., Kim, Y.-A., and Kim, J. H. (2019). Fight or flight? Crime as a driving force in business failure and business mobility. *Social Science Research*, 82:164–180.
- Hitsch, G. J., Hortacsu, A., and Lin, X. (2021). Prices and promotions in US retail markets. *Quantitative Marketing and Economics*, 19:289–368.
- Hollenbeck, B. (2017). The economic advantages of chain organization. *RAND Journal of Economics*, 48(4):1103–1135.
- Hollenbeck, B. and Giroldo, R. Z. (2022). Winning big: Scale and success in retail entrepreneurship. *Marketing Science*, 41(2):271–293.
- Hollenbeck, B. and Uetake, K. (2021). Taxation and market power in the legal marijuana industry. *RAND Journal of Economics*, 52:559–595.
- Jackson, O. and Tran, T. (2020). Larceny in the product market: A hidden tax? *Federal Reserve Bank of Boston Research Department Working Papers No. 20-14*.
- Janke, K., Propper, C., and Shields, M. A. (2016). Assaults, murders and walkers: The impact of violent crime on physical activity. *Journal of Health Economics*, 47:34–49.
- Janssen, A. and Zhang, X. (2023). Retail pharmacies and drug diversion during the opioid epidemic. *American Economic Review*, 113(1):1–33.
- Jia, P. (2008). What happens when wal-mart comes to town: An empirical analysis of the discount retailing industry. *Econometrica*, 76(6):1263–1316.
- Keiser, D. A. and Shapiro, J. S. (2019). Consequences of the clean water act and the demand for water quality. *Quarterly Journal of Economics*, 134:349–396.
- Klopach, B. (2024). One size fits all? The value of standardized retail chains. *RAND Journal of Economics*, 55(1):55–86.
- 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.

- Lewis, N. (2023). What the panic over shoplifting reveals about American crime policy. <https://www.themarshallproject.org/2023/02/27/shoplifting-retail-theft-lawmakers-response>. Retrieved June 7, 2023.
- Linden, L. and Rockoff, J. E. (2008). Estimates of the impact of crime risk on property values from Megan's laws. *American Economic Review*, 98(3):1103–1127.
- Lipscomb, M. and Mobarak, A. M. (2017). Decentralization and pollution spillovers: Evidence from the re-drawing of county borders in Brazil. *Review of Economic Studies*, 84:464–502.
- Lynch, A. K. and Rasmussen, D. W. (2001). Measuring the impact of crime on house prices. *Applied Economics*, 33(15):1981–1989.
- Mejia, D. and Restrepo, P. (2016). Crime and conspicuous consumption. *Journal of Public Economics*, 135:1–14.
- 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.
- Miravete, E. J., Seim, K., and Thurk, J. (2018). Market power and the Laffer curve. *Econometrica*, 86:1651–1687.
- Moore, R., Ricks, K., and Little, J. (2022). Annual State and Local Government Finances Summary: 2020. https://www2.census.gov/programs-surveys/gov-finances/tables/2020/2020_alfin_summary_brief.pdf. Retrieved March 10, 2024.
- 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.
- Nadreau, T. P., Fortenbery, T. R., and Mick, T. B. (2020). 2020 Contributions of the Washington cannabis sector. <https://app.leg.wa.gov/committeeschedules/Home/Document/234124#:~:text=In%202020%20the%20Washington%20cannabis,18%2C700%20full%20time%20equivalent%20jobs>. Retrieved March 9, 2022.
- National Retail Federation (2020). 2020 Organized retail crime survey. <https://nrf.com/research/2020-organized-retail-crime-survey>. Retrieved July 5, 2023.
- National Retail Federation (2022). Retail security survey. <https://cdn.nrf.com/sites/default/files/2022-09/National%20Retail%20Security%20Survey%20Organized%20Retail%20Crime%202022.pdf>. Retrieved November 17, 2023.
- National Retail Federation (2023). Retail security survey. https://cdn.nrf.com/sites/default/files/2023-09/NRF_National_Retail_Security_Survey_2023.pdf. Retrieved March 5, 2024.

- Owens, E. and Ba, B. (2021). The economics of policing and public safety. *Journal of Economic Perspectives*, 35(4):3–28.
- Parsa, H. G., Self, J. T., Njite, D., and King, T. (2005). Why restaurants fail. *Cornell Hotel and Restaurant Administration Quarterly*, 46:304–322.
- Phan, S. (2022). Pot retail stores consider armed security, other costly safety measures. <https://komonews.com/news/local/pot-retail-stores-consider-armed-security-proposed-safety-measures-will-be-costly>. Retrieved April 25, 2023.
- 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.
- Reagan, C. and Schlesinger, J. (2019). Inside home depot’s efforts to stop a growing theft problem at its stores. <https://www.cnbc.com/2019/11/22/inside-home-depots-efforts-to-stop-a-growing-theft-problem.html>. Retrieved June 8, 2023.
- 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.
- Ritz, R. A. (2024). Does competition increase pass-through? *RAND Journal of Economics*, 55(1):140–165.
- Rosenthal, S. S. and Ross, A. (2010). Violent crime, entrepreneurship, and cities. *Journal of Urban Economics*, 67(1):135–149.
- 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.
- Rozo, S. V. (2018). Is murder bad for business? Evidence from Colombia. *Review of Economics and Statistics*, 100:769–782.
- Saldanha, A. (2022). Armed robberies at WA pot shops hit decade high. <https://www.seattletimes.com/seattle-news/law-justice/armed-robberies-at-wa-pot-shops-hit-decade-high/>. Retrieved June 9, 2023.
- Schmidheiny, K. and Siegloch, S. (2023). On event studies and distributed-lags in two-way fixed effects models: Identification, equivalence, and generalization. *Journal of Applied Econometrics*, 38(5):695–713.

- Statista (2022). Average sales per store of U.S. supermarkets 2012-2016. <https://www.statista.com/statistics/240948/average-sales-per-store-of-us-supermarkets/>. Retrieved April 23, 2024.
- Statista (2024). In-store sales per convenience store in the United States from 2014 to 2022. <https://www.statista.com/statistics/308780/in-store-sales-per-convenience-store-in-the-us/>. Retrieved April 24, 2024.
- Stigler, G. J. (1970). The optimum enforcement of laws. *Journal of Political Economy*, 78(3):526–536.
- Stolkin, G. (2023). Paying for violence: The effect of drug trafficking organization presence on consumer prices. *Working Paper. Available at SSRN 4428016*.
- Stroebel, J. and Vavra, J. (2019). House prices, local demand, and retail prices. *Journal of Political Economy*, 127:1391–1436.
- Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199.
- Target (2023). Target closes select stores to prioritize team member and guest safety [Press Release]. <https://corporate.target.com/press/statement/2023/09/target-closes-select-stores-to-prioritize-team-member-and-guest-safety>. Retrieved February 8, 2024.
- Thornton, R. A. and Thompson, P. (2001). Learning from experience and learning from others: An exploration of learning and spillovers in wartime shipbuilding. *American Economic Review*, 91(5):1350–1368.
- U.S. Census Bureau (2022). U.S. Retail Sales Top \$6,523 Billion. <https://www.census.gov/newsroom/press-releases/2022/retail-sales.html>. Retrieved February 29, 2024.
- U.S. Chamber of Commerce (2022). Letter to the members of the United States congress. <https://www.uschamber.com/economy/us-chamber-letter-on-organized-retail-crime-and-retail-theft>. Retrieved October 3, 2023.
- U.S. Chamber of Commerce (2023). Crime risk to business: 2023 [Report]. https://www.uschamber.com/assets/documents/USCC-CrimeRisktoBusinessReport_FINAL.pdf. Retrieved February 21, 2024.
- 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>.

- Washington State Department of Financial Institutions (2022). Financial services options for cannabis retailers. <https://dfi.wa.gov/cannabis/financial-services>. Retrieved April 23, 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>. Retrieved April 24, 2024.
- Washington State Legislature (2015). Title 314. chapter 55.
- Washington State Legislature (2023). Bill Number SB 5259 - 2023-24.
- Washington State Liquor and Cannabis Board (2020). Bill Number E2SHB 2870.
- Washington State Liquor and Cannabis Board (2022). Board caucus meeting - cannabis retail safety forum. https://lcb.wa.gov/sites/default/files/publications/board/2022_Board_Agendas/03%2029%2022%20BOARD%20CAUCUS%20MINUTES%20-%20Signed.pdf. Retrieved June 15, 2023.
- Washington State Liquor and Cannabis Board (2024). Violations and due process. <https://lcb.wa.gov/enforcement/violations-and-due-process>. Retrieved April 10, 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.
- Washington State Office of the Attorney General (2022). Attorney General Ferguson convenes Washington Organized Retail Crime Theft Task Force [Press Release]. <https://www.atg.wa.gov/news/news-releases/attorney-general-ferguson-convenes-washington-organized-retail-crime-theft-task>.
- 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*.

Online Appendix

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


Figure A.1: Demographic characteristics of regular cannabis consumers




(A) Age




(B) Household income



(C) Education



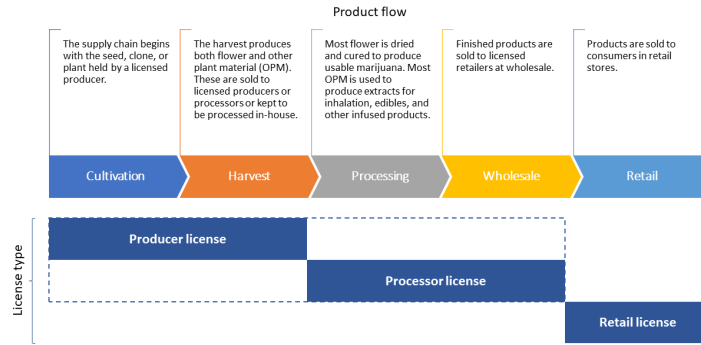
(D) Race/Ethnicity



(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.

Figure A.2: The cannabis supply chain

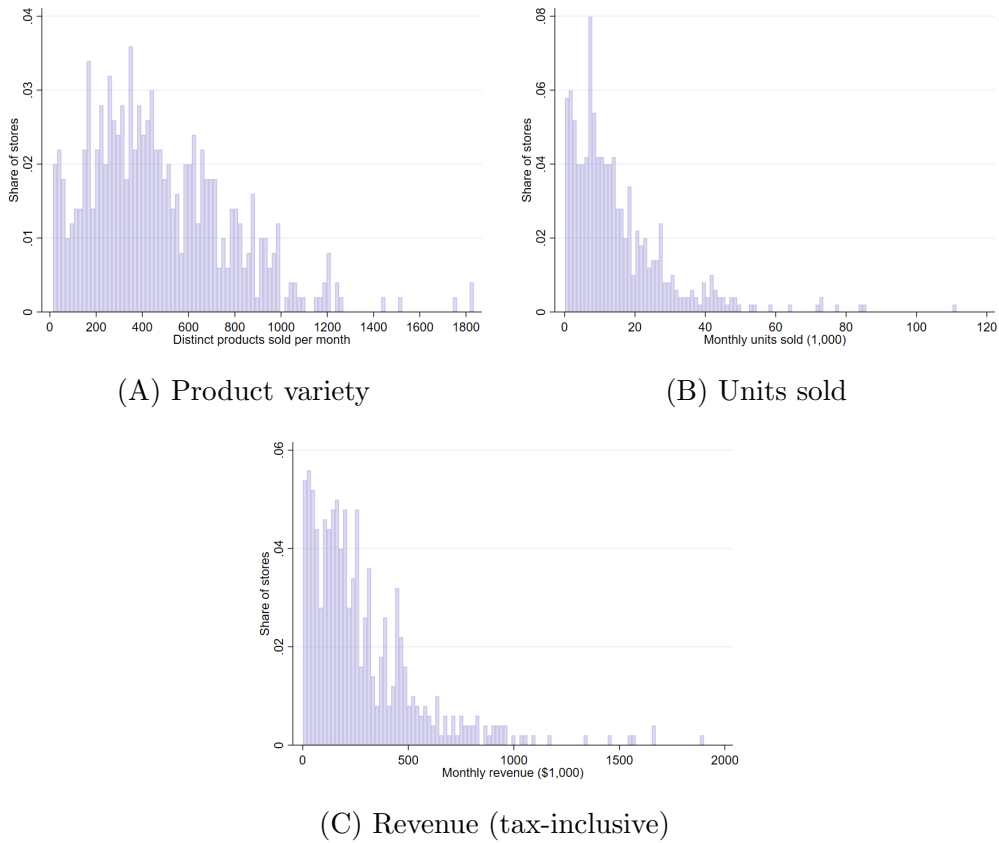


Notes: This figure depicts the flow of cannabis products, from left to right, as they move through the supply chain. Only licensed producers are permitted to cultivate and harvest cannabis plants; producers can only sell to licensed processors, who in turn are permitted to process products; only processors 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 473 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 14,836 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 \$285,320 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. 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
2018	\$324,582	\$702,358	0.32	0.68
2019	\$370,897	\$1,187,462	0.24	0.76
2020	\$407,273	\$1,584,301	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 2018-2020. Aggregate payroll data on cannabis retailers is from the Washington State Employment Security Department and High Peak Strategy (2018-2020). 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 5.22% 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 62% 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, 23.9% 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	5.22%	14.67%	15.59%	62.08%	23.90%	100%

Notes: This table shows the share of retailers’ wholesale expenditure according to wholesalers’ geographic proximity. The shares are based on 5.92 million unique wholesaler-retailer-product-month observations from August 2018 through July 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 Seasonality of crimes

Table 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

$$P_{i,j,t} = \frac{TR_{i,j,t}}{TQ_{i,j,t}}. \quad (8)$$

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 :

$$I_{c,j,t} = \prod_i \left(\frac{P_{i,j,t}}{P_{i,j,t-1}} \right)^{\omega_{i,c,j,y(t)}} \quad (9)$$

where each subcategory is a unique category-unit weight combination.³³ 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 .³⁴

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

$$I_{j,t} = \prod_c I_{c,j,t}^{\omega_{c,j,y(t)}}. \quad (10)$$

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

$$\pi_{j,t} = \ln I_{j,t} \quad (11)$$

B.2 Quantity indexes

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

$$I_{c,j,t} = \prod_i \left(\frac{Q_{i,j,t}}{Q_{i,j,t-1}} \right)^{\omega_{i,c,j,y(t)}} \quad (12)$$

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 wholesale cost indexes are constructed similar to the price indexes. TSD calculates the average monthly wholesale price at the SKU-level as

$$W_{i,j,t} = \frac{TE_{i,j,t}}{TQ_{i,j,t}}. \quad (13)$$

³³Since 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.

³⁴As 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.

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 :

$$I_{c,j,t} = \prod_i \left(\frac{W_{i,j,t}}{W_{i,j,t-1}} \right)^{\omega_{i,c,j,y(t)}} \quad (14)$$

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 . Note that by calculating this 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.³⁵

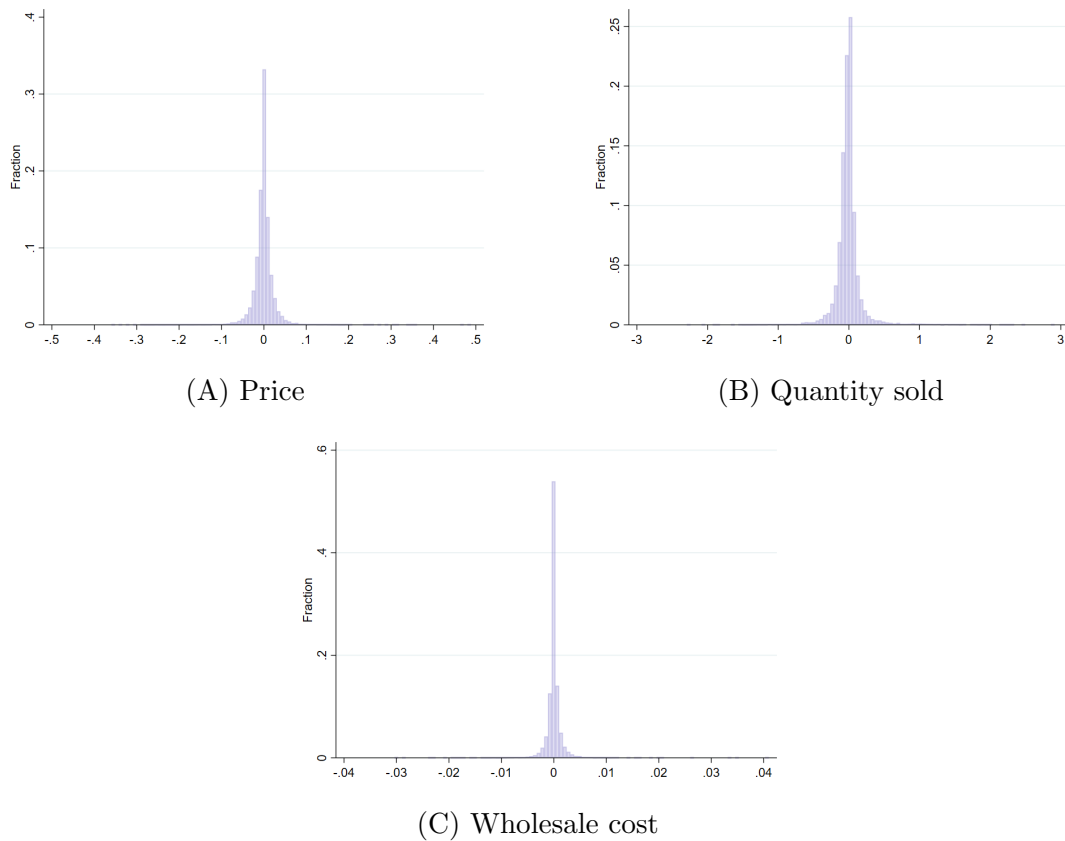
In the second step, we aggregate across subcategories to get the wholesale cost index for store j in month t :

$$I_{j,t} = \prod_c I_{c,j,t}^{\omega_{c,j,y(t)}}. \quad (15)$$

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, $\ln I_{j,t}$.

³⁵Results using monthly weights are available on request.

Figure A.5: Distributions for store-level indexes



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.3: 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.3 reports descriptive statistics for the different store types, based on all 511 active cannabis retailers.

In Table A.5 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.4 shows the descriptive statistics for our store-level market concentration measure, as defined in our main text.

Table A.4: Market concentration descriptive statistics

	(1) Mean	(2) Median	(3) Std
Victimized HHI	0.15	0.11	0.19
All HHI	0.27	0.15	0.28

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. Finally, we 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 0.007 while the difference in effect sizes in Section 5.2 is 0.022. 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 0.008 (Table 4, 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.5: Heterogeneity analysis robustness checks

	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
E_0	0.017 (0.012)	0.0014 (0.0032)	0.013 (0.010)	0.0017 (0.0027)	0.0013 (0.0044)	0.00097 (0.0019)	0.00066 (0.0025)	0.0018 (0.0031)
E_2	0.029** (0.014)	0.0060 (0.0051)	0.025** (0.012)	0.0067 (0.0051)	0.0025 (0.0050)	0.0052 (0.0037)	0.0019 (0.0037)	0.0081 (0.0053)
E_4	0.035** (0.016)	0.0074 (0.0076)	0.031** (0.014)	0.0068 (0.0074)	0.0094* (0.0053)	0.017*** (0.0062)	0.010* (0.0063)	0.021*** (0.0072)
E_6	0.023 (0.020)	0.023*** (0.0082)	0.026 (0.017)	0.019** (0.0086)	0.011* (0.0067)	0.011 (0.0072)	0.010 (0.0079)	0.013* (0.0073)
\sum Pre-event	0.010 (0.0089)	-0.0075 (0.0068)	-0.0055 (0.0092)	0.0044 (0.0062)	-0.00011 (0.011)	0.0040 (0.0064)	-0.0034 (0.0070)	0.0096 (0.0079)
N	11,968	18,793	15,004	15,757	28,889	29,904	16,023	15,001

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 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.

D.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:

$$\frac{\partial p_j}{\partial \tau} = \frac{\partial p_j}{\partial \alpha_j} \frac{\partial \alpha_j}{\partial \tau} + \sum_{i \neq j} \frac{\partial p_j}{\partial p_i} \frac{\partial p_i}{\partial \alpha_i} \frac{\partial \alpha_i}{\partial \tau} \quad (16)$$

where $\frac{\partial p_j}{\partial p_i}$ is firm j 's best response to a change in firm i 's price.³⁶ 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.

D.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.³⁷ 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).

³⁶For ease of exposition we consider competition in prices. Muehlegger and Sweeney (2022) show that this framework extends to a broad class of oligopolistic settings.

³⁷Wholesale 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.

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 model in first-differences:

$$\Delta p_{i,j,t} = \rho \Delta w_{i,j,t} + \sum_{r=1}^R \beta_r \Delta w_{i,r(j),t} + \Delta \gamma_t + \Delta \varepsilon_{i,j,t}, \quad (17)$$

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.

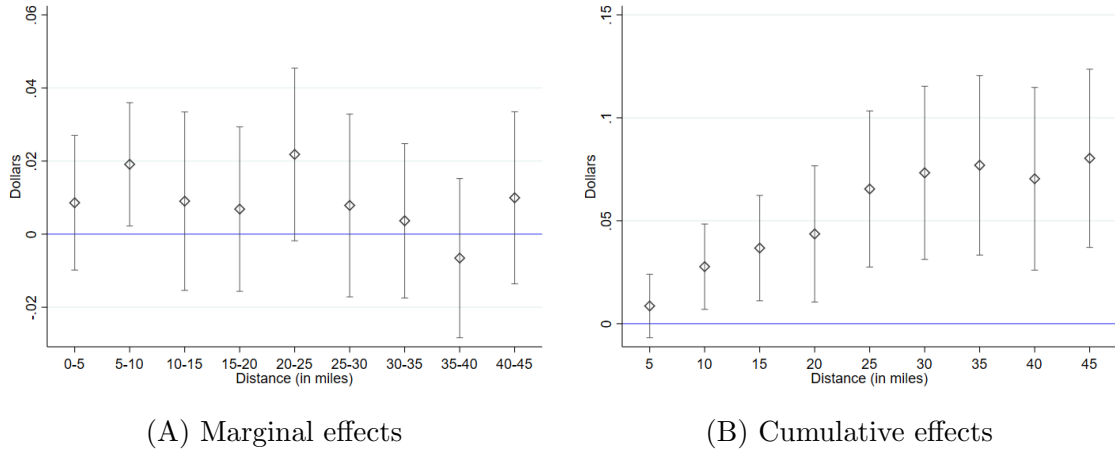
The effect of an aggregate change in unit costs on store j ’s prices comprises two parts. The first is the own cost pass-through rate, ρ , i.e. the increase in retail unit price at store j from the increase in store j ’s own wholesale unit cost. The second part is β_r which measures the pass-through 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.6 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.6. 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).

Figure A.6: 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 17). 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.

In addition to our main specification, we report several variants of our pass-through regression in Table A.6. In Column 2, we add controls (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). 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 17 in levels rather than first-differences, with store-product FE (Column 4). In Column 5, we specify equation 17 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.6 show that stores' sensitivity to competitors' costs plateaus at the 30-mile mark across 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 \$0.07 increase in retail prices solely due to

strategic complementarities. When estimated in levels (Column 4), prices increase \$0.48 from a \$1 increase in wholesale cost.

Table A.6: 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	1.735*** (0.0515)	1.735*** (0.0515)	1.735*** (0.0515)	1.217*** (0.426)	0.736*** (0.00991)	1.247*** (0.213)
Competitors' wholesale cost						
< 5 miles	0.00860 (0.00935)	0.00864 (0.00936)	0.00880 (0.00939)	0.0682* (0.0349)	0.000241 (0.00239)	-0.0167 (0.0617)
< 10 miles	0.0277** (0.0126)	0.0278** (0.0126)	0.0280** (0.0126)	0.231*** (0.0641)	0.00356 (0.00323)	-0.00656 (0.0844)
< 15 miles	0.0367** (0.0155)	0.0368** (0.0155)	0.0379** (0.0155)	0.250*** (0.0818)	0.00283 (0.00381)	-0.0224 (0.110)
< 20 miles	0.0436** (0.0200)	0.0437** (0.0201)	0.0445** (0.0201)	0.321*** (0.0946)	0.00618 (0.00479)	-0.0346 (0.135)
< 25 miles	0.0654*** (0.0230)	0.0656*** (0.0230)	0.0664*** (0.0230)	0.346*** (0.104)	0.00964* (0.00559)	-0.0615 (0.151)
< 30 miles	0.0733*** (0.0255)	0.0733*** (0.0255)	0.0734*** (0.0255)	0.481*** (0.129)	0.0162*** (0.00611)	-0.0581 (0.167)
< 35 miles	0.0769*** (0.0264)	0.0768*** (0.0264)	0.0766*** (0.0264)	0.511*** (0.144)	0.0157*** (0.00597)	-0.0852 (0.187)
< 40 miles	0.0704*** (0.0269)	0.0705*** (0.0269)	0.0702*** (0.0269)	0.491*** (0.149)	0.0152*** (0.00532)	-0.0291 (0.214)
< 45 miles	0.0803*** (0.0262)	0.0805*** (0.0263)	0.0799*** (0.0262)	0.446*** (0.145)	0.0203** (0.00856)	0.00962 (0.229)
<i>N</i>	2,290,818	2,290,726	2,290,818	3,683,684	2,290,818	9,729

Notes: The table reports the estimates of wholesale cost pass-through rates from equation 7. 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. Columns 5-6 are interpretable as pass-through elasticities. Standard errors are clustered at the store level and shown in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Overall, the results from Figure A.6 and Table A.6 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.6 and Table A.6 as providing supportive evidence for our definition of unaffected local markets from Section 4.

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

It is worth noting that equation 17 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.6. The own-cost pass-through elasticity is 1.25 aligns with our baseline finding that costs are more than fully passed through. However, the coefficients on competitors' costs fluctuate and are not statistically significant, in line with the idea that competitors' cost indexes may not capture strategic complementarity in prices for the reasons outlined above.

E Spatial treatment and alternative definitions of unaffected local markets and rival stores


This section elaborates on the inclusion criteria for our control group, with a special emphasis on the geographical criteria. Subsequent subsections conduct several robustness checks to assess the sensitivity of our estimated effects to these geographical criteria and our definitions of rival stores.

E.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.7. 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.7: 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.7, 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.7. 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 .³⁸ 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.7, 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.

³⁸Note 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.

E.2 Alternative inner and outer ring sizes

In our main specification, we set $\bar{d}^r = 5$, $\underline{d}^c = 30$ and $\bar{d}^c = 60$. To investigate whether our estimated effects hinge on our definition of clean control stores, we estimate our model using alternative spatial definitions of clean control and rival stores. Figures A.8(A) and A.8(B) show the cumulative treatment effects four months after a crime incident for victimized stores and for rival stores, respectively. The middle bar refers to our main specification. The first bar in both figures shows the effects when expanding the inner boundary (\underline{d}^c) to include all stores within 10 and 60 miles as potential control stores, and the third bar shows the effects of expanding the outer boundary (\bar{d}^c) to include all stores that are located more than 30 miles away from victimized stores.

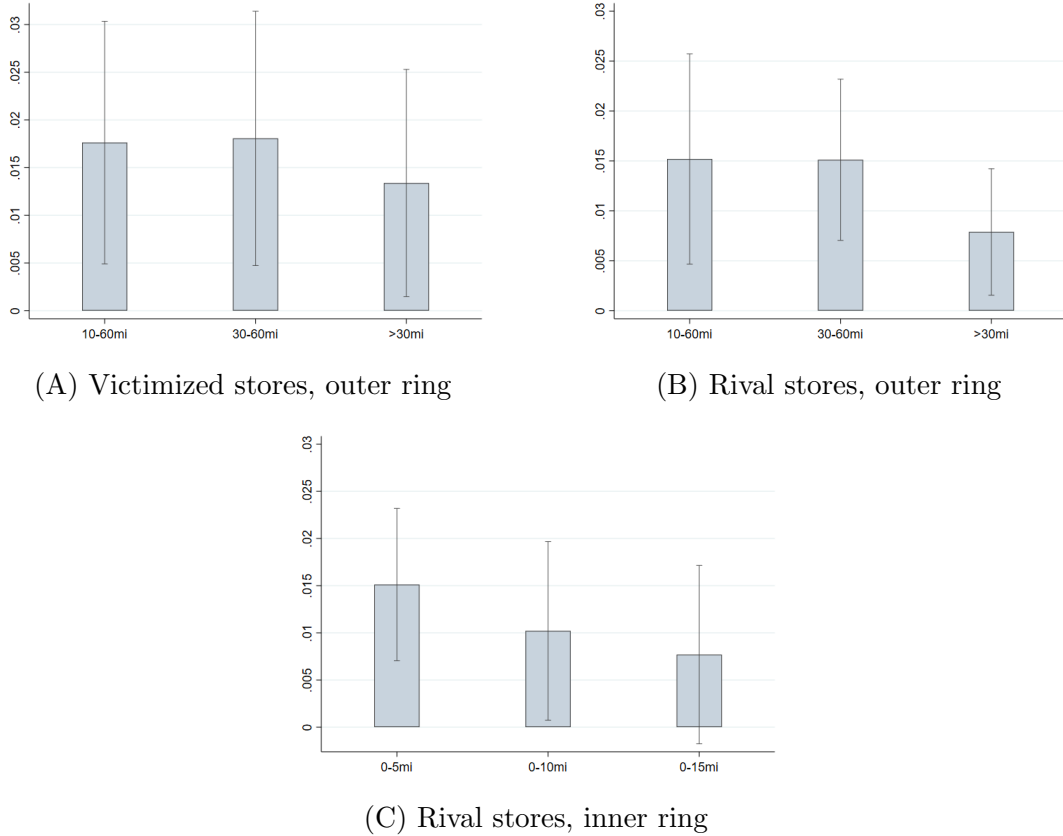
Similarly, Figure A.8(C) illustrates estimated effects as we expand the definition of the inner ring defining rival stores (\bar{d}^r). The first bar reflects cumulative effects in our main specification; the second shows effects for stores within 10 miles, and the third for stores within 15 miles of victimized stores.

E.3 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 two 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 E.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.9 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. Using data from the Washington State Office of Financial Management, 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

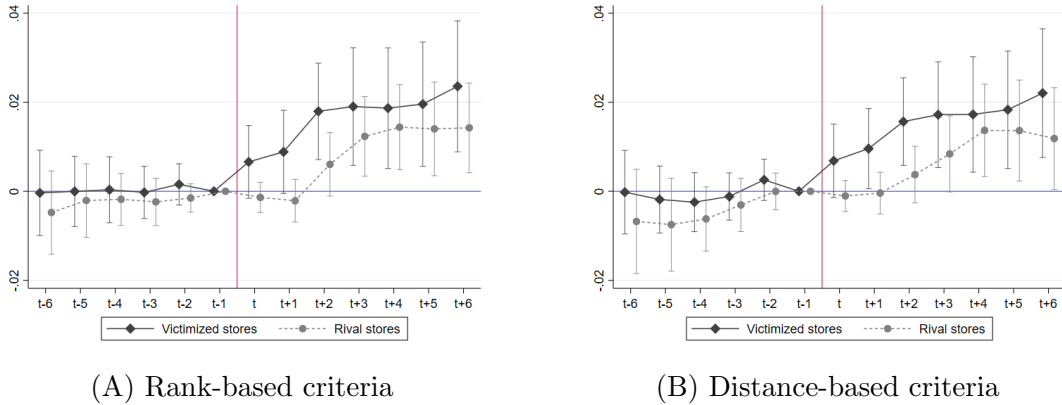
Figure A.8: Price effects with alternative inner and outer rings



Notes: Panels A and B show the cumulative treatment effects four months after a crime incident for victimized stores and for rival stores, respectively, along with corresponding 90% confidence intervals based on standard errors clustered at the store level. The first bar in both panels shows the effects when considering all stores within 10 and 60 miles as potential control stores. The middle bar refers to our main specification. The third bar shows the effects when considering all stores located more than 30 miles away from victimized stores as clean control stores. Panel C illustrates estimated effects as we expand the definition for rival stores. The first bar reflects cumulative effects in our main specification; the second bar shows effects when considering all stores within 10 miles as rival stores, and the third bar shows effects when considering all stores within 15 miles as rivals.

of rival stores remains 5 miles (results are similar if we reduce the radius for urban rivals). Panel B of Figure A.9 shows results for this specification. Estimated treatment effects for victimized stores 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. The price effect at rival stores is 1%, a slight attenuation compared to the baseline specification. Nevertheless, the timing and path of treatment effects are identical to the baseline specification. Taken together, the results from Figure A.9 show that our main results hold when we account for geospatial heterogeneities across urban and rural areas.

Figure A.9: Price effects with spatial heterogeneity



Notes: Panels A and B 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.

E.4 Treatment spillovers from other crime incidents

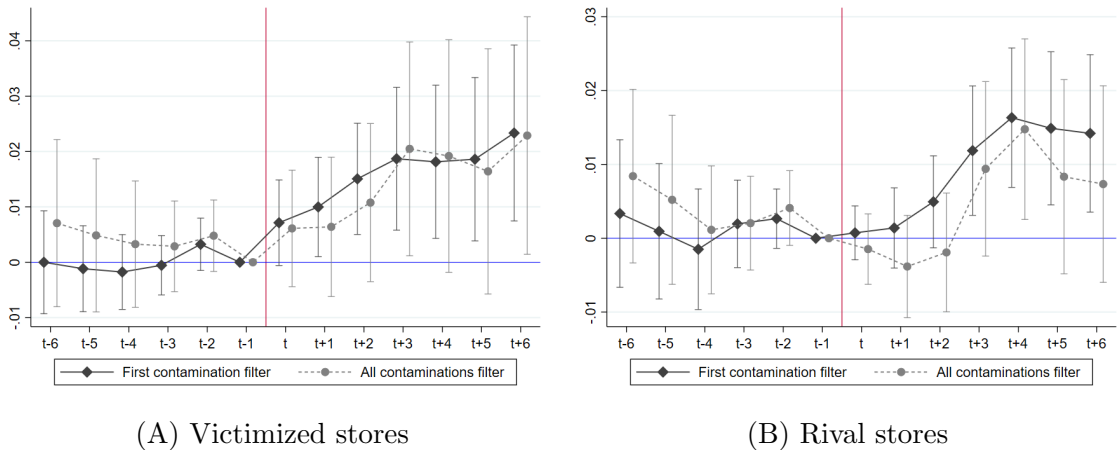
As discussed above, a store must be at least 30 miles from a victimized store to qualify as a clean control in a sub-experiment. However, unlike victimized and rival stores, this distance restriction does not apply to stores between \bar{d}^r and \underline{d}^c from other victimized stores with overlapping 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.7, 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:

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

Figure A.10 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.10). 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.10: 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.

E.5 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 242 to 225 and three sub-experiments drop out of the estimation sample. However, Figure A.11 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 1.3% 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.11: 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.

F 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.³⁹ 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 329 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.⁴⁰ 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 aver-

³⁹Supporting this line of argument, including rivals in the victimized store specification slightly decreases the treatment effects, as expected, that remain statistically significant at the 5% level for all alternative specifications four months after treatment (results not shown).

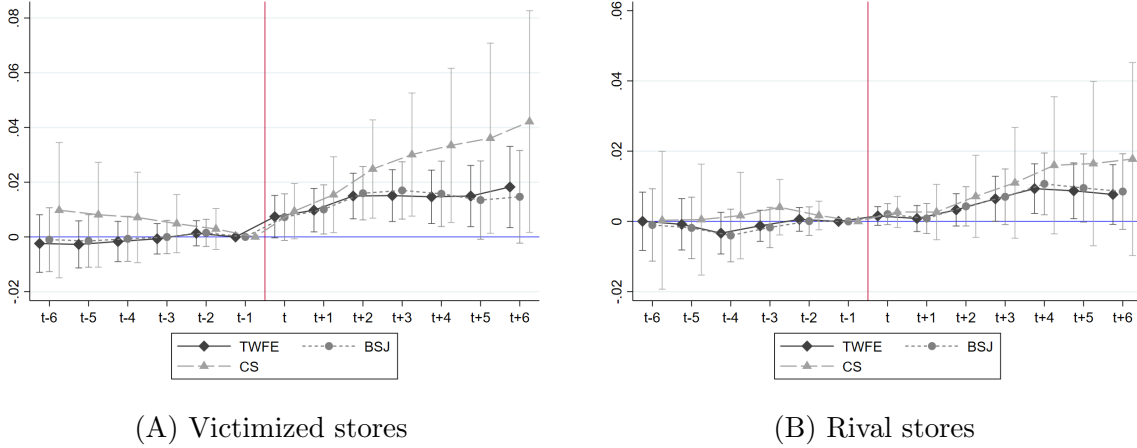
⁴⁰Some observations (one victimized store and a few rival stores) are dropped through the "autosample" option. We cannot impute non-treated potential outcomes for those stores, implying that the BSJ estimator cannot estimate unbiased treatment effects for these stores.

age 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 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.12 presents the cumulative treatment effects for the alternative estimators. For victimized stores, the effect size three months after the incident is approximately 1.5% using the TWFE estimator and 1.7% with the BSJ estimator (Panel A, Figure A.12). Both coefficients are statistically significant at the 1% 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 3.0% higher three 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%. 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.12: 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).

G Additional robustness checks

G.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.13. 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.

G.2 Longer event window

Figure A.14 illustrates that our main results carry over when we extend the event window from 12 months to 18 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.13: 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 12-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 little effect on wholesale prices, as in the main section. Taken together, the results from Figure A.14 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.


G.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.15. Considering the store-level quantity index has a standard deviation ten times larger than the price index and 100 times larger than the wholesale cost index (Table 2), the quantity index estimates may be particularly sensitive to outliers. Figure A.15 Panel A shows that winsorizing decreases the standard errors of our rival treatment estimates. Notably, winsorizing our data also mitigates the slight pre-trend for the quantity sold at rival


Figure A.14: Effect of crime on store outcomes with extended event window



(A) Price level



(B) Quantity sold

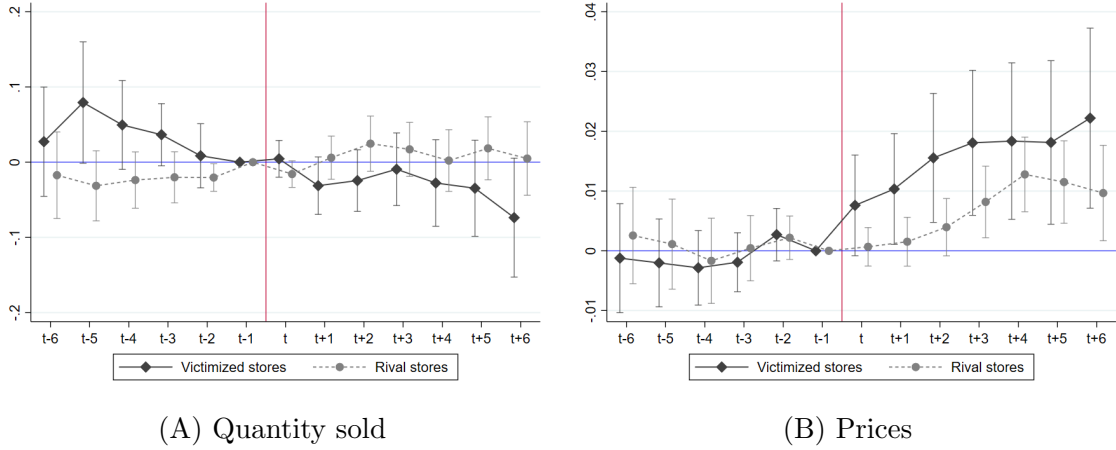


(C) Wholesale cost

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).

stores while leaving the effects at victimized stores unchanged. This suggests that the slightly significant pre-trend in the quantity specification for rival stores in Figure 3 is primarily due to a few outliers. Panel B, corresponding to the estimates in Column 4 of Table 5, indicates that winsorizing the store-level price indexes does not influence the estimated price effects of retail crime, further reaffirming our main results.

Figure A.15: Effect of 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%.

G.4 Store closures

One concern is that the costs associated with retail crime incidents 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 57 victimized stores in our main specification, eight drop out before the end of the 46-month sample period, a dropout rate of 0.14. For these eight stores, the average duration between a crime incident and dropping out is 15 months. Only three stores drop out within the first 12 months after an incident, which is equal to 5.3% of victimized stores. For comparison, 94 of the 450 non-victimized stores drop out during the sample period, a rate of 0.21, or 5.5% on average per year. For comparison, the annual dropout rate for restaurants is estimated to be about 30% (Parsa et al., 2005). Taken together, we view this as suggestive evidence that crime incidents do not lead to a higher dropout rate.