The Price of Violence: The Impact of Social Unrest on Real Estate in Santiago, Chile.

Ignacio Aravena-Gonzalez*

July 27, 2026

Click here for the latest version

Abstract

This paper estimates the effect of urban violence on real estate markets using the October 2019 social unrest in Santiago, Chile. I combine geocoded violence exposure from prosecutorial records with the universe of residential and commercial property transactions and a large panel of rental listings in a difference-in-differences design with block, year-month, and municipality-by-year fixed effects. Severely exposed blocks trade at a persistent residential discount of about 5.6% up to six years after the shock, while blocks with low exposure show no effects; applied to the exposed housing stock, this implies a capitalized loss of roughly USD 1.6–1.7 billion. Asking rents fall during the disruption and then recover, so prices and rents diverge; the implied capitalization rate on exposed property rises by about 30 basis points. Residential transaction activity does not contract and if anything increases, indicating rotation rather than illiquidity. Commercial markets adjust on the opposite margin: prices are uninformative, but the probability of any transaction falls by roughly 36%. Exposed blocks also lose about 3% of their population. A channel decomposition and triple-difference design attributes these effects to violence rather than to metro-accessibility disruption, as blocks near closed stations without violence show no discount, and greater accessibility disruption is associated with smaller violence discounts. Some mechanisms that help to explain these effects are based on market thickness and changes in amenities (security) and commercial contraction. The discount concentrates in central, mid and higher SES blocks with pre-existing insecurity. The results suggest that the economic cost of urban violence depends not only on its intensity, but also on the margin through which each market adjusts.

Keywords: Urban violence, social unrest, house prices, capitalization rates, market thickness.
JEL Codes: R21, R31, R33, K42, O18.

*London School of Economics and Political Science. i.a.aravena-gonzalez@lse.ac.uk

1 Introduction

In recent decades, major cities around the world have repeatedly experienced episodes of mass civil unrest. From Hong Kong and Paris to Beirut, Bogotá, and Santiago, protests have at times escalated into sustained periods of disorder and violence that disrupted transport systems, damaged commercial corridors, and altered the functioning of urban neighborhoods (Haig et al., 2020). These episodes raise an economic question on when concentrated urban violence persists long enough to disrupt how a city functions, does it leave durable scars on local property markets, or do values and market activity recover once the unrest subsides? A shock of this magnitude can damage physical structures, disrupt local commerce, and change perceptions of safety and neighborhood reputation, thereby revising forward-looking beliefs about the viability and value of a location well beyond the period of violence itself.

Yet even when violence lowers the value of a place, that loss might not become visible in the same way across real estate markets. In relatively active markets, such as residential, it is more likely to be capitalized into observed prices as transactions continue and valuations adjust. In thinner markets, such as commercial, the response may instead take the form of lower participation, delayed transactions, or a more selected set of observed sales, so that prices capture only part of the underlying economic cost. Households and firms do not adjust in the same way, and commercial assets depend more directly on expected rental income, tenant demand, and local business conditions. Hence, the same episode of urban violence may produce residential repricing but commercial market thinning, even when the underlying local shock is common. A related distinction arises within the residential market itself, since sale prices capitalize the entire expected future stream of housing services, while rents reflect current willingness to pay, so the may not move together, and the gap between them is informative about what exactly changed.

A limited economic literature shows that violent shocks can leave persistent scars on residential demand, property values, and neighborhood trajectories, though existing estimates often rely on city-level aggregates, highly aggregated spatial units, small treatment groups, or coarse temporal variation, which limits what can be learned about localized adjustment within cities (Collins and Margo, 2007; Besley and Mueller, 2012; Abadie and Dermisi, 2008). Related work also suggests that hedonic responses are larger when crime signals personal risk to residents rather than mere property loss (Gibbons, 2004; Linden and Rockoff, 2008), but this evidence focuses almost entirely on residential sale prices, and rarely on prices and rents jointly. Evidence on commercial property is thinner still; Lens and Meltzer (2016) document negative effects of crime on commercial values in New York City using assessed values rather than market transactions, but that evidence does not reveal whether commercial adjustment operates through prices, participation, or both. More broadly, we still know little about whether the same episode of urban violence affects residential and commercial real estate through similar mechanisms, or whether one market adjusts primarily through repricing while the other adjusts through market thinning.

This paper studies the October 2019 social unrest in Santiago, Chile, the largest episode of civil unrest in the country's recent democratic history. What began as protests against a small increase in metro fares quickly escalated into a broader wave of demonstrations, public disorder, arson, and looting. Prosecutorial records linked to the unrest document 32,901 cases corresponding to 35,146 criminal incidents (Fiscalia Nacional, 2025), while official metro reports indicate that 118 of the network's 136 stations suffered some form of damage during the first days of the crisis (Metro de Santiago, 2020). Santiago therefore offers a rich source of within-city variation in violence exposure, with administrative records detailed enough to trace both prices and market activity at a fine spatial scale (block levels). At the same time, this variation is not randomly distributed. Violence was concentrated around metro infrastructure, central locations, and commercially active corridors, and the unrest was followed first by a prolonged metro disruption and then by the pandemic. The empirical challenge is to determine whether the persistent effects observed in exposed areas reflect violence itself or these overlapping shocks.

To address this challenge, I combine geocoded violence exposure from the national prosecutor's office with the universe of residential and commercial property transactions, a large panel of rental listings, firm-level tax records, and census data. The empirical strategy compares outcomes before and after October 2019 across census blocks with different levels of violence exposure, absorbing block fixed effects, year-month fixed effects, and municipality-specific annual shocks. Rather than treating the post period as homogeneous, I distinguish the immediate unrest period from a later post-normalization phase beginning in October 2022, when metro service had resumed and pandemic-related distortions were no longer the dominant explanation for local real estate dynamics. I also distinguish between low-exposure blocks, which intersect a single incident type, and more severely exposed blocks, which intersect multiple incident types; furthermore, I exploit the imperfect overlap between violence hotspots and damaged metro stations to separate violence-related revaluation from direct accessibility disruption. Blocks near closed stations but without violence exposure exhibit no persistent discount and in fact trade at a small premium, while severely exposed blocks far from damaged stations show the full effect, making a purely accessibility-based interpretation difficult to sustain. A complementary triple-difference design using continuous GTFS-based network disruption sharpens this: the violence discount is largest in blocks with no accessibility loss at all, and among exposed blocks those with greater network disruption exhibit smaller discounts, which is the opposite of what a pure accessibility channel would predict.

I then study both prices and market participation, asking whether the effects of violence appear through repricing or through market thinning. In residential real estate, the response appears in prices. Severe violence exposure is associated with a persistent sale-price discount of about 5.6% in the post-normalization period, three to six years after the shock, which in a back-of-the-envelope calculation implies a capitalized loss of roughly 42 million UF across the severely exposed residential stock, or between USD 1.6 and 1.7 billion. Participation does not fall; if anything severely exposed blocks transact 4.5% more often, which points to rotation rather than illiquidity and implies that the discount is realized rather than merely posted. In commercial real estate the pattern inverts, as sale-price effects are imprecisely estimated and statistically indistinguishable from zero, while the probability that a block-quarter records any transaction falls by 0.78 percentage points, roughly a 36% reduction relative to pre-shock control-block baselines. Administrative firm records corroborate this, as incumbent businesses do not exit the registry, but the number of establishments reporting at least one worker falls by about 13% in activities dependent on street-level demand, with no comparable effect in a placebo group of office-based activities. Local commerce thins out on the employment margin while remaining administratively alive.

Two further results move beyond prices. First, asking rents in severely exposed blocks fall during the disruption and the pandemic and then recover most of that decline, while sale prices continue to fall. Prices and rents therefore diverge, with an implied gross capitalization rate on exposed residential property rises by about 4.8 log points, roughly 30 basis points on a baseline yield of 6.2%, and the estimate is stable across specification families and spatial units. This wedge is consistent with a higher required return on locations perceived as riskier and equally with lower expected growth in local rents–the two are not separately identified–. Second, the revaluation has a counterpart as severely exposed blocks lose about 3% of their population between the 2017 and 2024 censuses. Both results support reading the discount as a change in the expected future value of these locations rather than as a transitory dislocation of the transaction market.

A simple conceptual framework, developed in Section 2, helps interpret these patterns. Violence can reduce the expected surplus from location in both residential and commercial real estate, but the margin through which that loss becomes visible depends on market structure. In relatively liquid markets with active matching, shocks are more readily capitalized into observed prices (Wheaton, 1990; Krainer, 2001). In thinner markets, the same shock may instead reduce participation, making observed transaction prices a more selected and less complete signal of underlying value (Fisher et al., 2003; Rosenthal and Ross, 2010). This is especially relevant in commercial real estate, where firms can respond to local violence by postponing entry, downsizing, relocating, or simply not transacting in affected areas. In that case, null observed price effects need not imply zero economic cost; they may instead reflect adjustment through non-trading. The framework also implies that if violence changes the risk or the expected growth of the income stream rather than only its current level, sale prices and rents should separate, which is what the cap-rate result documents. Finally, it suggests that more severe and visible exposure should have stronger effects than isolated or less salient incidents, consistent with distinguishing low from severe exposure in the empirical design.

These patterns are consistent with a common shock that becomes visible through different margins across markets. In residential real estate the response appears in sale prices, in the price-rent wedge, and in a net loss of residents, while transaction activity remains active. In commercial real estate it appears through reduced participation and through the thinning of employers, with observed prices carrying almost no information. The contrast is anticipated by the raw coverage of the two markets, since 87% of severely exposed blocks record a residential transaction over the sample period, against only 24% for commercial.

The paper contributes to four strands of literature. First, it adds to the literature on conflict, riots, crime, and property markets by providing evidence from a large urban episode using geocoded administrative microdata at census-block scale, a level of spatial resolution that existing studies, which typically rely on city-level aggregates, highly aggregated spatial units, or small treatment groups, have not been able to exploit (Collins and Margo, 2007; Besley and Mueller, 2012; Abadie and Dermisi, 2008). Second, it contributes to the literature on crime capitalization by showing that sufficiently severe urban violence generates persistent residential effects, with little corresponding response in low-exposure areas, and by documenting that these effects are more consistent with violence than with temporary metro-accessibility disruption (Gibbons, 2004; Linden and Rockoff, 2008; Gourley, 2019). Third, it contributes to the emerging literature on commercial property and urban violence by moving from assessed values to market transactions and showing that the main commercial response is lower participation rather than clean repricing (Rosenthal and Ross, 2010; Fisher et al., 2003). Fourth, by estimating price and rent responses jointly within the same episode, it shows that a disamenity shock need not be summarized by a single discount: prices and rents move by different amounts, and the implied capitalization rate is itself an outcome, one that a literature focused on sale prices alone cannot observe.

The remainder of the paper proceeds as follows. Section 2 describes the October 2019 unrest and institutional setting. Section 3 presents the data and exposure measures. Section 4 outlines the empirical strategy. Section 5 reports the main results on prices, rents, capitalization rates, transaction activity, and population. Section 6 examines mechanisms, including market thickness, firm-level evidence, and changes in reported crime. Section 7 reports heterogeneity. Section 8 presents robustness checks. Section 9 concludes.

2 Institutional Setup and Conceptual Framework

2.1 The October 2019 Social Unrest

On October 18, 2019, protests against a CLP 30 increase[1] in Santiago's metro fare escalated into the largest episode of civil unrest in Chile's recent democratic history (Gonzalez and Morán, 2020; Cox et al., 2024). What began as student-led fare evasion quickly expanded into broader demonstrations over inequality, pensions, healthcare, and the cost of living, among other grievances (Avendaño, 2019; Sasse, 2019). Within days, protests spread across Santiago and other major cities, with large marches centered on Plaza Italia, later renamed "Plaza de la Dignidad" by protesters, which became the symbolic heart of the movement (Gonzalez and Morán, 2020).

Figure 1 traces the evolution of protest activity using the Conflict and Social Cohesion Ob-

[1] CLP 30 is roughly USD 0.033 and GBP 0.025. The increase was lower than 5%.

servatory (COES) event database (COES, 2020). Protest and public-disorder events rose sharply after October 18 and remained elevated well into early 2020, with particularly high intensity in the Metropolitan Region. These data are useful for fixing the timing and scale of the unrest, even though the block-level exposure measure used below comes from geocoded prosecutorial records rather than protest-event counts.

The protests were accompanied by substantial violence and property damage. An official statistical report released by the Chilean Public Prosecutor's Office in October 2025 consolidates all criminal cases linked to the unrest between October 18, 2019 and March 31, 2020 (Fiscalia Nacional, 2025). The report records 32,901 cases corresponding to 35,146 criminal incidents, including looting, public-disorder offenses, arson, violence against state agents, and other crimes associated with the unrest. These records provide the basis for the block-level violence exposure measure used throughout the paper.

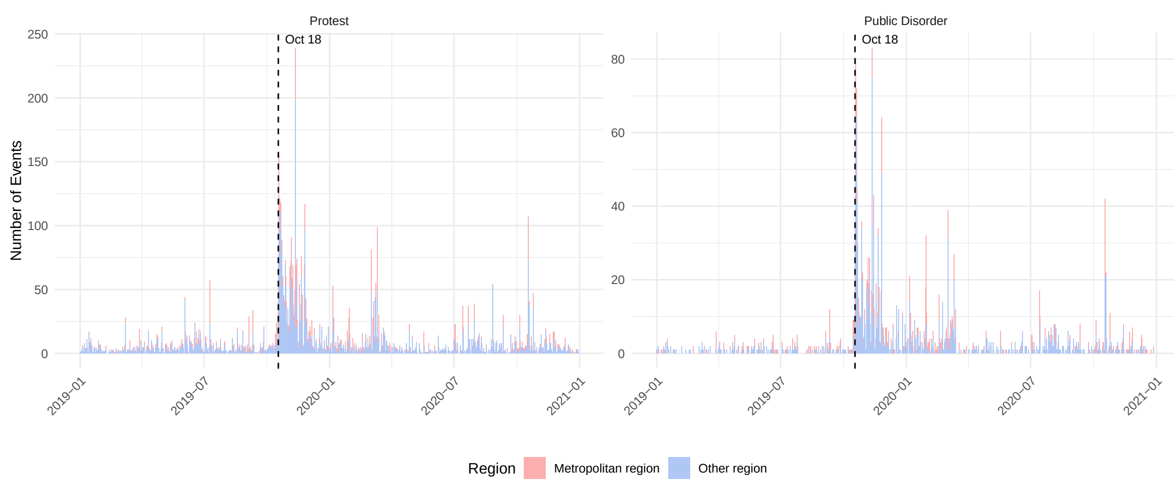

Figure 1: Protest intensity over time. Weekly counts of protest events and public-disorder events from the COES database [3]. The vertical dashed line marks October 18, 2019.

2.2 Violence, Metro Disruption, and the Pandemic

The October 2019 unrest in Santiago is empirically challenging because the protests overlapped with a major transit disruption and was followed soon after by the pandemic. Localized violence varied sharply across census blocks, while damage to the metro network was widespread, as 118 of the system's 136 stations suffered some form of damage, 25 were affected by fire, and full network operations were restored only on September 25, 2020 (Metro de Santiago, 2020). The pandemic then introduced an additional citywide shock, altering commuting patterns, retail activity, and the value of central accessibility.

These shocks were intertwined in both time and space. Violence was concentrated along major corridors and around metro infrastructure, so blocks more exposed to unrest were also more likely to be close to damaged stations. At the same time, the deterioration observed in affected areas cannot be reduced to transit disruption alone. Contemporary market reporting described sharp commercial deterioration in and around the city's epicenter, including rent reductions, store closures, and rising vacancy in the areas most affected by repeated unrest.[2]

This is why the analysis distinguishes four phases. The pre-shock period runs up to October 2019. The metro-closure period spans October 2019 to September 2020, when violent unrest and severe transport disruption coincided. The COVID-transition period runs from October 2020 to September 2022, when metro service had resumed but commuting patterns and urban demand remained distorted by the pandemic and the shift toward remote work. The post-normalization period begins in October 2022, when neither the acute transit shock nor the pandemic transition remains the leading explanation for local real estate dynamics. A persistent discount after that point is therefore harder to attribute to temporary station closures alone.

Figure 2 summarizes the spatial footprint of the unrest. Violence was not confined to a single enclave. It extended across multiple parts of Santiago and cut across neighborhoods with different socioeconomic profiles, providing meaningful within-city variation in exposure and making it less likely that the results are driven by a single urban submarket.

2.3 Conceptual Framework: Prices, Participation, and Market Thickness

The question is not only whether urban unrest lowers local real estate values, but also through which margin that loss becomes visible in the data. A localized episode of violence can reduce the expected surplus from location in several related ways. It may worsen perceived safety, increase visible disorder, damage the street environment, or weaken local commercial conditions, all of

[2]As descriptive market evidence rather than causal identification, *Diario Financiero* reported rent declines of up to 50% in commercial areas around Plaza Italia in 2020 (Diario Financiero, n.d.). Later reporting based on GPS Property documented roughly 900 commercial closures in Santiago Centro and office vacancy of 14.9% by 2024 (EMOL, n.d.). Even by the end of 2025, CBRE still measured retail vacancy in Santiago Centro at 9.28% (La Tercera, n.d.). These sources are useful as descriptive evidence of local market deterioration, but they do not substitute for the transaction-based analysis developed below.

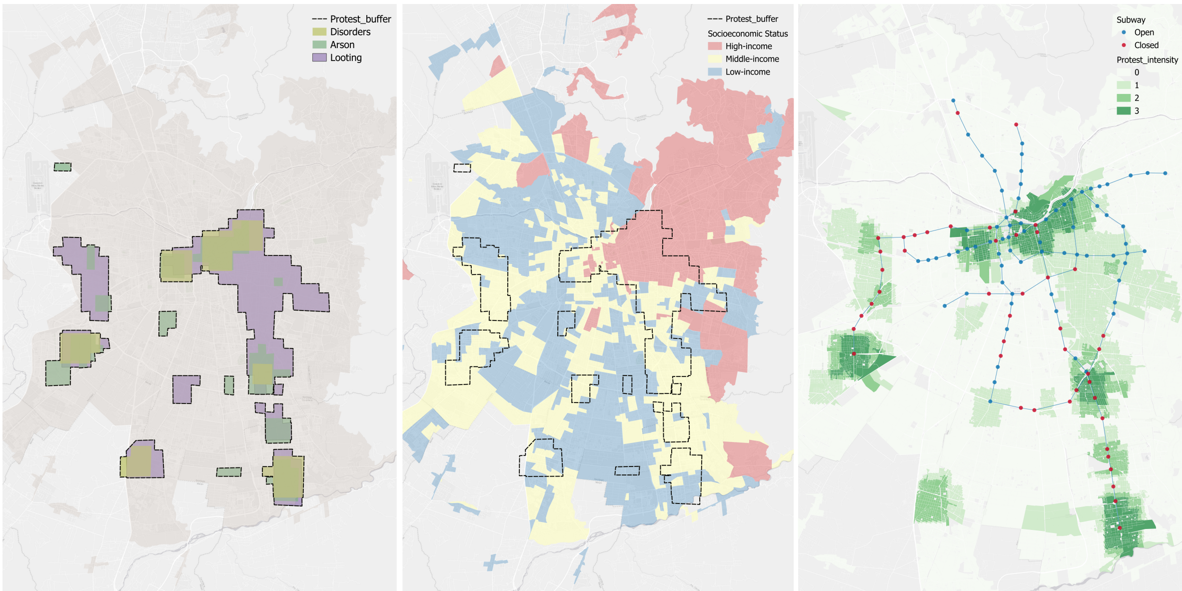

Figure 2: Spatial distribution of violence in Santiago. The left panel shows geocoded criminal incidents linked to the unrest. The center panel overlays neighborhood socioeconomic status with violence buffers. The right panel shows the violence intensity index used in the empirical analysis together with the metro network and damaged stations [4,5].

which can affect how households and firms value a place (Ooi and Le, 2013; Zahirovich-Herbert and Gibler, 2014). In Santiago, these effects overlapped with temporary metro disruption, so the empirical objective is not to recover a perfectly isolated violence channel. It is instead to assess whether the persistent discount observed in exposed areas is more consistent with localized unrest than with a temporary accessibility shock.

The same shock need not appear in the same way across markets. Prices are observed only when transactions occur. In relatively active markets, shocks are more readily capitalized into observed prices. In thinner markets, the same shock may instead reduce the likelihood of trade, lengthen search, and make observed transactions a selected subset of the underlying stock. I use the term *market thickness* to refer to this combination of transaction frequency, liquidity, and ease of matching in local real estate markets (Wheaton, 1990; Krainer, 2001). This distinction matters especially in commercial real estate, where values depend on expected rental income and business profitability, and where a negative local shock may reduce not only the price conditional on trade, but also the incentive to lease, enter, or transact at all (Fisher et al., 2003; Rosenthal and Ross, 2010).

This helps organize the residential and commercial results. In residential markets, household mobility is often slower and adjustment at the extensive margin is more limited, in part because moving and search are costly (Hanushek and Quigley, 1979). A deterioration in neighborhood conditions may therefore be reflected more readily in lower sale prices, and potentially in lower asking rents, while transaction activity remains comparatively stable. In commercial markets, by contrast, the extensive margin is more flexible. Firms can postpone entry, relocate, delay leasing decisions, or simply avoid transacting in affected areas. Under that pattern, the main empirical response may be lower market participation rather than a clean repricing of observed sales. Null observed price effects therefore need not imply zero economic cost; they may instead indicate that the margin of adjustment is non-trading.

The framework also motivates distinguishing between low and severe exposure. Not all incidents convey the same signal about local conditions. More severe and visible forms of disorder are more likely to reinforce perceptions of local deterioration than isolated or less salient incidents (Sampson and Raudenbush, 1999; Bordalo et al., 2010). In the empirical design, this distinction is captured through exposure to multiple types of violent incidents rather than through raw incident counts alone. This provides a basis for allowing the response to differ across exposure categories, rather than imposing a single linear treatment gradient from the outset.

The empirical analysis is organized around three predictions. First, if violence rather than transit disruption is the main persistent force, severely exposed blocks should continue to display discounts even after metro service has normalized, including in locations not directly adjacent to damaged stations. Second, in the residential market, the response should appear primarily in transaction prices rather than in a persistent contraction in transaction activity, with rents providing a complementary outcome that may move in the same direction if local valuations decline. Third, in the commercial market, the same shock may appear primarily through lower transaction probability and market thinning. Under that pattern, null price effects do not imply zero economic cost; they imply that the cost is revealed through participation rather than through observed repricing alone.

3 Data

3.1 Data Sources

The analysis combines six data sources. First, I use administrative records of residential and commercial property transactions in the Santiago Metropolitan Region from January 2016 through October 2025 from the *Servicio de Impuestos Internos* (SII), Chile's tax authority. Each record reports the transaction price (in UF, an inflation-indexed unit of account), constructed area, lot size, year built, property type, and geographic location at the census-block (*manzana*) level.[3] I restrict the sample to transactions with non-missing price and key characteristics and trim prices at the 1st and 99th percentiles. The final residential sample contains 619,710 transactions; the commercial sample contains 31,023 transactions.

Second, I complement transactions with residential rental listings from TOCTOC, a major Chilean real estate portal. The listings report asking rents and unit attributes matched to census blocks. Including rents (43,044 listings) provides an independent outcome not subject to the same selection concerns as transaction prices: if violence reduces the valuation of a location, both sale prices and asking rents should respond.

Third, I use georeferenced incident data from the National Public Prosecutor's Office (Fiscalia Nacional, 2025), which consolidated criminal cases linked to the unrest filed between October 18, 2019 and March 31, 2020. The Prosecutor's Office provides statistically significant spatial clusters (*hotspots*) for three incident categories—public disorder, arson, and looting—validated through formal spatial analysis. I intersect these hotspot polygons with census blocks to assign exposure at the block level.

Fourth, I use operational records from Metro de Santiago identifying stations damaged during the unrest and the duration of closures (Metro de Santiago, 2020). I compute the distance from each block to the nearest closed station and define baseline metro exposure using a 500-meter threshold.[4]

Fifth, I supplement the analysis with police crime statistics from Carabineros de Chile's *Plan Cuadrante*, which provides annual crime counts by type at the police-quadrant level for 2018 and

[3] The UF (*Unidad de Fomento*) is updated daily to reflect CPI changes and is the standard unit for real estate pricing in Chile.

[4] Robustness checks in Section 8 consider thresholds from 300 to 800 meters.

2023. Sixth, I incorporate pre-treatment Census (2017) information on population and density at the block level and the ISMT socioeconomic index (Araya, 2023).

3.2 Exposure Measures

Discrete violence index. For each block, I define three binary indicators for intersection with a Prosecutor hotspot of (i) public disorder, (ii) arson, and (iii) looting, and sum them to obtain $V_b \in \{0, 1, 2, 3\}$. The index captures the variety of incident types to which a block was exposed rather than event counts within any category. Of the 48,803 blocks in the sample, 76% have $V_b = 0$, 15% have $V_b = 1$, and 9% have $V_b \geq 2$. Table 1 reports the composition of each level: 78% of $V_b = 1$ blocks are exposed to looting alone, while $V_b = 2$ blocks predominantly combine arson with looting (54%) or disorder with looting (39%). $V_b = 3$ blocks have all three types present by construction. Figure 3(a) maps the spatial distribution.[5]

Continuous intensity. I also construct a continuous measure $\tilde{V}_b$ by summing the rescaled Prosecutor heatmap values for each incident type at the block level, capturing within-category variation in exposure intensity (Figure 3(b)). The correlation with the discrete index is high (Spearman $\rho = 0.97$), but the continuous measure may offer additional variation within exposure categories. Both measures are used in the empirical analysis; their relative strengths are discussed alongside the results.

Metro disruption. I define $M_b = \mathbb{1}\{\text{distance to nearest closed station} \leq 500\text{m}\}$ and combine it with violence exposure to classify blocks into six groups crossing the three violence levels ($V_b = 0$, $V_b = 1$, $V_b \geq 2$) with metro proximity ($M_b = 0$, $M_b = 1$). This classification, shown in Figure 3(c), is used to separate violence effects from accessibility disruptions while preserving the three-level structure of the baseline specification. Additionally, I construct a network-based accessibility measure using GTFS transit data. For each block, I compute the shortest-path travel

[5] To assess whether the final hotspot geography was already spatially anchored at the outset of the unrest, I construct a continuous early-exposure measure using a kernel-density surface based on georeferenced protest points observed on October 18–19, 2019. Figure A.1 maps this early protest surface, while Appendix Table A.4 reports its association with the final block-level violence classification. The relationship is clearly positive in raw comparisons and after conditioning on baseline centrality, metro proximity, density, and local housing and commercial stock, and it remains positive and statistically significant even after introducing municipality fixed effects. I therefore interpret this exercise as descriptive evidence that the broad geography of unrest was already visible at the outset, rather than as proof that final severe-hotspot status was fully predetermined at the block level.

time to the CBD under both pre-closure and post-closure station configurations and take the difference as the potential accessibility shock (Figure 3(d)). Among affected blocks, the mean shock is 25.5 minutes and the median is 19.9 minutes. Quintiles of baseline accessibility, interacted with year indicators, serve as flexible controls for differential trends across accessibility levels.

As Figure 3 makes clear, violence exposure, metro disruption, and accessibility shocks are spatially correlated: protests deliberately targeted metro infrastructure, so the geography of damaged stations overlaps substantially with the geography of violent incidents. This overlap means that a simple comparison of exposed and unexposed blocks would confound violence-related revaluation with transit-related accessibility changes. The channel classification and the GTFS-based controls are designed to address this, but the underlying spatial correlation limits the precision with which individual channels can be separately identified—a point I return to in the mechanism analysis of Section 7.

Table 1: Composition of Violence Exposure Index

V_b	Combination	Blocks	% of V_b	Transactions
$V_b = 0$: *No exposure*				
	(none)	37,066	100.0	394,713
$V_b = 1$: *One incident type*				
	Looting only		78.4	
	Arson only		17.1	
	Public disorder only		4.5	
	Subtotal	7,522		125,903
$V_b = 2$: *Two incident types*				
	Arson + Looting		53.8	
	Disorder + Looting		39.4	
	Disorder + Arson		6.8	
	Subtotal	2,528		43,041
$V_b = 3$: *All three types*				
	Disorder + Arson + Looting	1,687	100.0	56,053
Total exposed ($V_b > 0$)		11,737		224,997
Total		48,803		619,710

Notes: Percentages show the share of each combination within the violence level, computed on blocks with at least one residential transaction. Block totals and transaction counts from the full sample (48,803 blocks, 619,710 residential transactions).

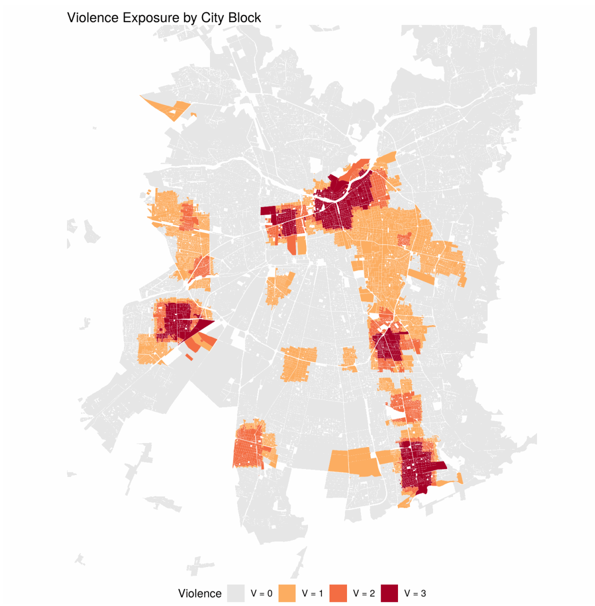

(a) Discrete violence index V_b

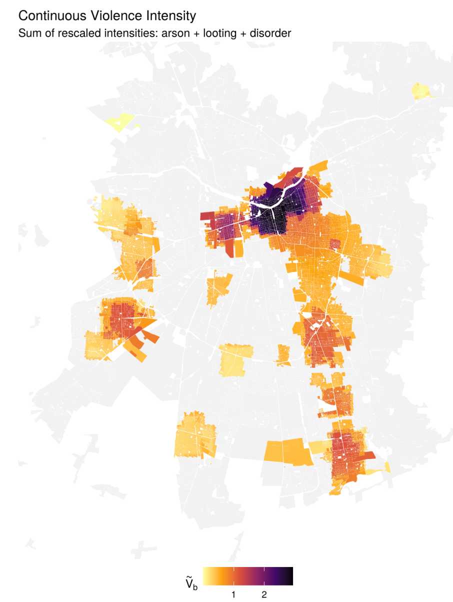

(b) Continuous intensity $\tilde{V}_b$

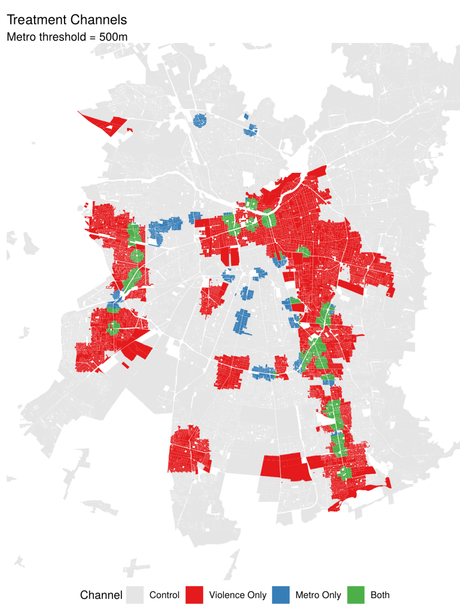

(c) Treatment channels

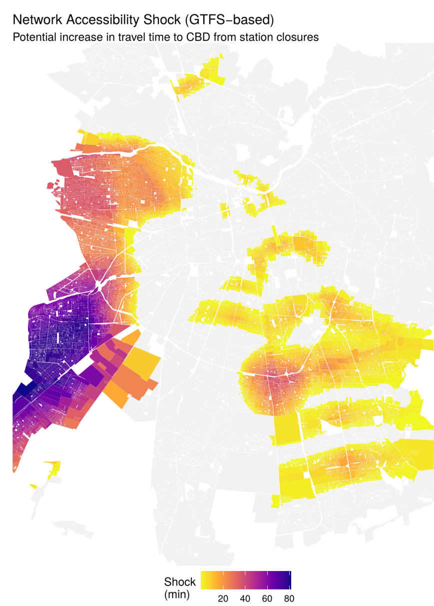

(d) Network accessibility shock (GTFS)

Figure 3: Geography of Violence Exposure and Metro Disruption. Panel (a) maps the discrete violence index $V_b \in \{0, 1, 2, 3\}$ based on Prosecutor hotspot intersections. Panel (b) maps the continuous intensity $\tilde{V}_b$, summing rescaled heatmap values for arson, looting, and public disorder. Panel (c) classifies blocks into treatment channels using $V_b \geq 2$ and proximity to closed metro stations ($\leq$ 500m). Panel (d) maps the GTFS-based potential accessibility shock in minutes of additional travel time to the CBD from station closures.

3.3 Descriptive Statistics

Table 2 reports pre-period summary statistics by violence exposure. Exposed blocks are more central, closer to metro stations, and have higher commercial presence, reflecting the non-random geography of the unrest. The gradient steepens with exposure: $V_b = 3$ blocks average 654 m from the nearest metro station and contain 18.7 commercial properties per block, compared with 6,183 m and 1.4 in control areas. Residential prices follow the same gradient, increasing from 43.5 UF/m^2 in control areas to 55.8 UF/m^2 in $V_b = 3$. This pattern is not explained by socioeconomic status alone—the GSE level increases only modestly from 1.6 to 2.0 across exposure groups—but rather reflects the concentration of violence around transit nodes and commercially active corridors. Identification in the difference-in-differences design relies on differential changes over time, conditional on block and flexible location-time controls, rather than on balance in pre-period levels.

The last column of Table 2 reports characteristics for the pooled $V_b \geq 2$ group. This group has an apartment share of 85%, intermediate between $V_b = 2$ (78%) and $V_b = 3$ (90%), and an average metro distance of 1,276 m. The compositional gradient between $V_b = 2$ and $V_b = 3$ is steeper than between $V_b = 1$ and $V_b = 2$, a point discussed further in Section 8.

Figure 4 plots hedonic price indices by violence intensity—residuals from a regression of log price on property characteristics and block and year-month fixed effects, normalized to 100 in October 2017. The series track closely before October 2019. After the shock, higher-intensity areas display gradually widening declines relative to the control group, with $V_b = 2$ and $V_b = 3$ separating more visibly than $V_b = 1$. The divergence is gradual rather than immediate, consistent with slow capitalization rather than a short-lived disruption.[6]

Figure 5 reports event-study estimates for selected residential covariates, testing whether observable property characteristics evolve differentially in violence-exposed blocks ($V_b > 0$) relative to controls. The coefficients are generally small and statistically insignificant in both the pre- and post-shock periods. Joint F-tests for post-shock stability fail to reject at the 5% level for three of four covariates; the apartment share shows marginal significance ($p = 0.045$), consistent with a modest compositional shift that is explicitly controlled for in robustness specifications including

[6] Analogous indices for commercial transactions and rental listings are reported in Appendix Figure A.2 and Figure A.3.

Table 2: Summary Statistics by Violence Exposure (Pre-Period)

	V=0 (Control)	V=1 (Low)	V=2 (Medium)	V=3 (High)	V≥2 (Pooled)
Panel A: Residential transactions					
Price (UF/m^2)	43.5	49.2	53.1	55.8	54.6
	(26.1)				
Total price (UF)	3,113	3,060	3,213	3,016	3,103
	(3,529)				
Construction (m^2)	66.1	62.2	61.0	56.0	58.2
	(39.1)				
Year built	2003	2002	2000	2000	2000
Apartment (=1)	0.64	0.76	0.78	0.90	0.85
Observations	176,330	53,126	18,193	22,878	41,071
Panel B: Commercial transactions					
Price (UF/m^2)	63.5	59.5	58.6	59.4	59.2
	(40.8)				
Total price (UF)	7,161	9,731	8,952	6,857	7,470
Construction (m^2)	162	206	183	151	161
Year built	1997	1990	1986	1985	1985
Office (=1)	0.62	0.68	0.62	0.69	0.67
Observations	6,775	2,461	1,573	3,800	5,373
Panel C: Block characteristics					
Residential stock	40.9	51.0	54.3	95.7	70.9
Commercial stock	1.4	2.5	5.9	18.7	11.0
Distance to metro (m)	6,183	1,586	1,690	654	1,276
Pop. density (per ha)	198	209	180	194	185
Population 2017	137	147	142	198	165
GSE level	1.6	1.7	1.8	2.0	1.9
Blocks	37,066	7,522	2,528	1,687	4,215
Panel D: Rental listings					
Rent (UF/m^2)	0.30	0.30	0.30	0.30	0.30
	(0.10)				
Rent (UF)	28.7	21.5	22.0	17.3	19.3
	(15.7)				
Size (m^2)	96.5	75.5	76.2	57.6	65.5
Bedrooms	2.3	2.1	2.0	1.7	1.8
Bathrooms	2.1	1.8	1.8	1.5	1.6
Observations	4,741	1,845	917	1,258	2,175

Notes: Standard deviations in parentheses (shown for the control group). Pre-period: before October 2019. $V_b \in \{0, 1, 2, 3\}$ based on Prosecutor hotspot intersections. The last column pools $V_b = 2$ and $V_b = 3$. Full sample sizes: 619,710 residential transactions, 31,023 commercial, 43,044 rental listings across 48,803 blocks.

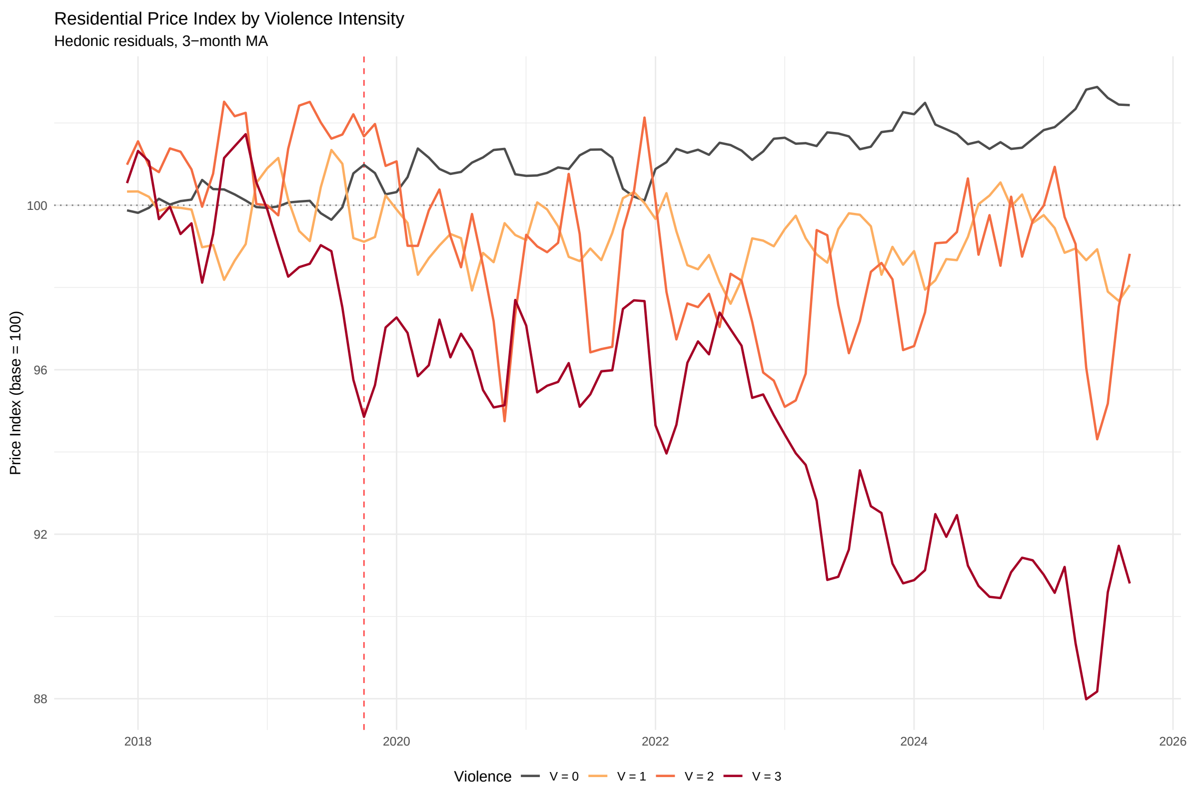

Figure 4: Residential Price Indices by Violence Intensity. Hedonic residuals from a regression of log price on property controls with block and year-month fixed effects, normalized to 100 in October 2017 and smoothed using 3-month moving averages. The vertical dashed line marks October 2019.

property type × year fixed effects.[7]

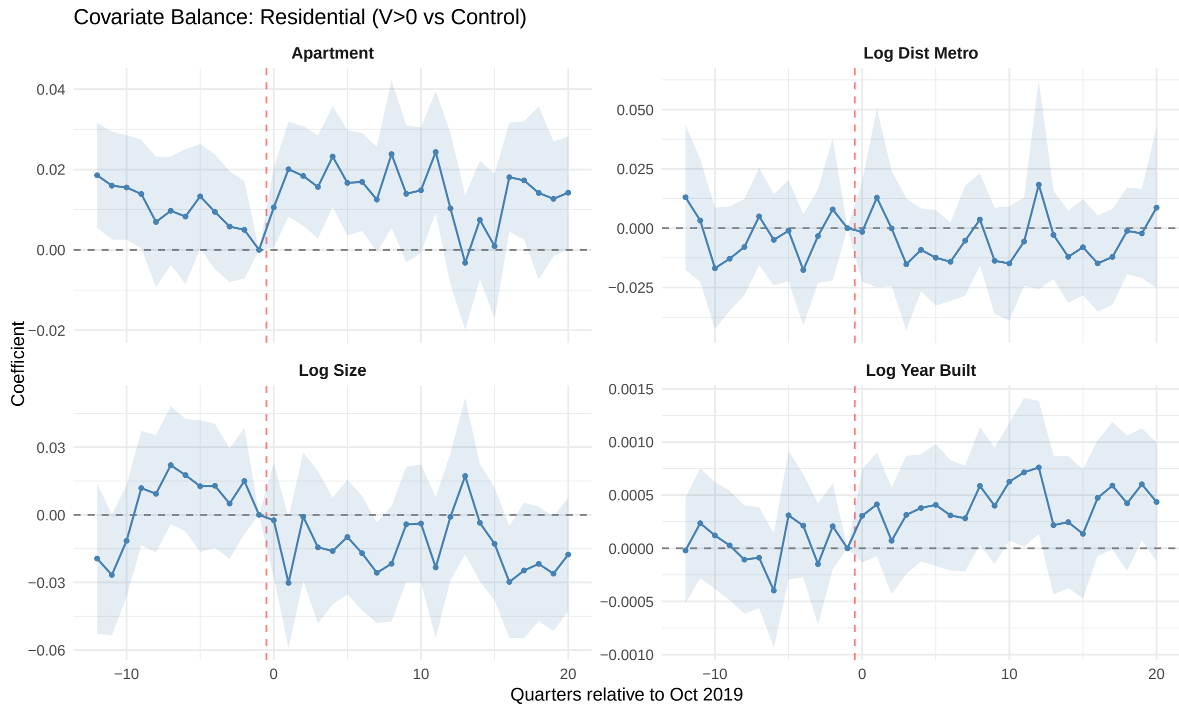

Figure 5: Covariate Balance Event Studies (Residential). Coefficients from regressions of each covariate on $\mathbb{1}\{V_b > 0\}\times$ quarter indicators with block and year-month fixed effects. Bands: 95% CIs, SEs clustered at block level. Reference: $Q{-}1$. Joint F-tests for post-shock coefficients: log size $p = 0.071$, log year built $p = 0.508$, apartment $p = 0.045$, log distance to metro $p = 0.425$.

4 Empirical Strategy

The October 2019 social unrest constitutes a sharp temporal shock whose spatial incidence varies across census blocks. Identification comes from comparing outcomes before and after the shock across blocks with different levels of exposure to unrest-related violence. The key identifying assumption is that, absent the unrest, outcomes in more- and less-exposed blocks would have evolved similarly after conditioning on block fixed effects, time fixed effects, and municipality-specific annual shocks. Treatment onset is simultaneous—variation comes from cross-sectional differences in exposure severity and channel, not from staggered adoption.

[7]Commercial covariate balance (Appendix Figure A.4) shows no significant imbalance on any dimension (joint F: log size $p = 0.317$, log year built $p = 0.364$, office $p = 0.978$, log distance to metro $p = 0.858$). Rental listings (Appendix Figure A.5) show marginal imbalance for bedrooms ($p = 0.017$), consistent with the compositional gradient toward smaller units in exposed areas; other covariates pass (log size $p = 0.116$, bathrooms $p = 0.484$, log distance to metro $p = 0.811$).

A central threat to interpretation is that violence exposure is spatially correlated with centrality, commercial intensity, and metro infrastructure, which may themselves generate differential dynamics—particularly after the onset of COVID-19 and the shift toward remote work. I address this in two ways: first, by augmenting the baseline with accessibility-quintile-by-year fixed effects that flexibly absorb differential trends across accessibility levels; and second, by exploiting the institutional shock to the metro system to separate violence exposure from accessibility disruptions.

4.1 Baseline Difference-in-Differences

Let P_{ibt} be the transaction price of property i in block b at time t. The baseline specification is:

$$\begin{aligned}\ln P_{ibt} = \sum_{p} \Big[\beta_p^1 \left(\mathbb{1}\{V_b = 1\} \times \mathbb{1}\{t \in p\}\right) + \beta_p^{2+} \left(\mathbb{1}\{V_b \geq 2\} \times \mathbb{1}\{t \in p\}\right)\Big] \\ + X_{it}'\gamma + \alpha_b + \theta_t + \mu_{c(b),y(t)} + \varepsilon_{ibt},\end{aligned} \tag{1}$$

where p indexes four periods: *pre-shock* (before October 2019, the omitted category), *metro closure* (October 2019–September 2020), *COVID transition* (October 2020–September 2022), and *post-normalization* (October 2022 onward). Blocks with no violence exposure ($V_b = 0$) serve as the reference group. X_{it} includes property-level controls: log constructed area, log lot size, log year built, and a property-type indicator.[8] α_b are block fixed effects, θ_t are year-month fixed effects, and $\mu_{c(b),y(t)}$ are municipality-by-year fixed effects. Standard errors are clustered at the block level throughout.

The specification distinguishes low exposure ($V_b = 1$) from severe exposure ($V_b \geq 2$), with each group receiving its own set of period interactions relative to unexposed blocks. This distinction reflects the composition of the violence index. As documented in Table 1, $V_b = 1$ blocks are predominantly exposed to a single incident type—78% to looting alone—while $V_b = 2$ and $V_b = 3$ blocks are exposed to combinations of arson, public disorder, and looting. I pool $V_b = 2$ and $V_b = 3$ because both involve multiple incident types, they are fewer in number (2,528 and 1,687 blocks respectively, compared to 7,522 for $V_b = 1$), and distinguishing between two and three overlapping incident types is unlikely to identify meaningfully different treatment intensities.

To address the concern that persistent effects may reflect differential trends in transport-

[8] Distance to the nearest metro station is fixed at the block level and absorbed by α_b.

accessible areas rather than violence per se, I augment the baseline with accessibility-quintile-by-year fixed effects:

$$\begin{aligned} \ln P_{ibt} = & \sum_{p} \Big[\beta_p^1 \left(\mathbb{1}\{V_b = 1\} \times \mathbb{1}\{t \in p\} \right) + \beta_p^{2+} \left(\mathbb{1}\{V_b \geq 2\} \times \mathbb{1}\{t \in p\} \right) \Big] \\ & + X_{it}'\gamma + \alpha_b + \theta_t + \mu_{c(b),y(t)} + \sum_q \sum_y \phi_{qy} \left(\mathbb{1}\{Q_b = q\} \times \mathbb{1}\{y(t) = y\} \right) + \varepsilon_{ibt}, \end{aligned} \tag{2}$$

where $Q_b \in \{1, \ldots, 5\}$ denotes the quintile of pre-shock GTFS-based travel time to the CBD for block b. The ϕ_{qy} terms absorb any year-specific shock that differentially affects areas by accessibility level, including WFH-related revaluations. If the baseline violence effect were driven by accessibility trends, one would expect it to attenuate once these controls are included.

I estimate equations (1) and (2) for residential sale prices, residential asking rents, and commercial sale prices. As mentioned before, when sale prices and asking rents move in the same direction, that pattern is more suggestive of a change in location valuation than of composition effects alone, though the identifying content of the rental evidence is weaker than for transaction prices.

I also apply the same treatment structure to block-quarter transaction activity outcomes. For each block-quarter, I estimate a linear probability model for whether any transaction occurs (extensive margin) and Poisson pseudo-maximum likelihood (PPML) estimates for transaction counts.[9] This is particularly relevant in the commercial market, where transactions are sparse and negative local shocks may affect the likelihood of trade as well as the prices observed among transacting properties.

As a complementary specification, I replace the categorical treatment with the standardized continuous intensity $\tilde{V}_b^{std}$, which captures within-category variation in exposure. Both the categorical and continuous measures are reported in the main results.

4.2 Violence versus Accessibility

I implement two complementary designs to separate violence effects from metro-accessibility disruptions.

[9] PPML is estimated on the full panel including zeros and includes block fixed effects, quarter fixed effects, and municipality-by-year fixed effects analogous to the baseline specification.

Channel decomposition. The first approach estimates separate coefficients for six exposure groups, preserving the three-level violence classification from the baseline:

$$\begin{aligned}\ln P_{ibt} = \sum_p \Big[&\beta_p^1\big(\mathbb{1}\{V_b{=}1, M_b{=}0\}\cdot\mathbb{1}\{t\in p\}\big) + \beta_p^{1M}\big(\mathbb{1}\{V_b{=}1, M_b{=}1\}\cdot\mathbb{1}\{t\in p\}\big)\\ &+ \delta_p^{MO}\big(\mathbb{1}\{V_b{=}0, M_b{=}1\}\cdot\mathbb{1}\{t\in p\}\big) + \beta_p^{2+}\big(\mathbb{1}\{V_b{\geq}2, M_b{=}0\}\cdot\mathbb{1}\{t\in p\}\big)\\ &+ \beta_p^{B}\big(\mathbb{1}\{V_b{\geq}2, M_b{=}1\}\cdot\mathbb{1}\{t\in p\}\big)\Big] + X_{it}'\gamma + \alpha_b + \theta_t + \mu_{c(b),y(t)} + \varepsilon_{ibt},\end{aligned} \tag{3}$$

with unexposed blocks far from closed stations ($V_b = 0$, $M_b = 0$) as the omitted category. This preserves the distinction between unexposed, low-exposure, and severely exposed blocks established in the baseline, while allowing each group to interact separately with metro proximity. In the main text I focus on three key contrasts: (i) δ_p^{MO}, the effect of metro disruption absent any violence; (ii) β_p^{2+}, the effect of severe violence absent metro disruption; and (iii) β_p^{B}, the effect when both are present. If persistent price effects are driven by violence rather than accessibility, one would expect β_p^{2+} to be negative while δ_p^{MO} should be close to zero.

Triple-difference. The second approach uses the full sample and decomposes the treatment into violence levels, metro proximity, and their interactions:

$$\begin{aligned}\ln P_{ibt} = \sum_p \Big[&\beta_p^1\big(\mathbb{1}\{V_b{=}1\}\times\mathbb{1}\{t\in p\}\big) + \beta_p^{2+}\big(\mathbb{1}\{V_b{\geq}2\}\times\mathbb{1}\{t\in p\}\big) + \delta_p\big(M_b\times\mathbb{1}\{t\in p\}\big)\\ &+ \kappa_p^1\big(\mathbb{1}\{V_b{=}1\}\times M_b\times\mathbb{1}\{t\in p\}\big) + \kappa_p^{2+}\big(\mathbb{1}\{V_b{\geq}2\}\times M_b\times\mathbb{1}\{t\in p\}\big)\Big]\\ &+ X_{it}'\gamma + \alpha_b + \theta_t + \mu_{c(b),y(t)} + \varepsilon_{ibt}.\end{aligned} \tag{4}$$

Here β_p^1 and β_p^{2+} capture the effect of low and severe violence far from closed stations, δ_p captures the effect of metro proximity for unexposed blocks, and κ_p^1 and κ_p^{2+} test whether metro proximity modifies the violence effect at each level. This is the triple-difference analogue of the baseline: both violence levels are allowed to interact with metro disruption, rather than forcing low-exposure blocks into either the treatment or control side of the decomposition.[10]

Both designs are estimated for residential prices and rents. Commercial results are reported in the Appendix.

[10]I also estimate versions replacing M_b with a continuous measure of station-access disruption and versions interacting violence with GTFS-based travel time to the CBD. These are reported in the Appendix.

Pre-trends are assessed using event-time specifications at semiannual and quarterly frequencies, with joint Wald tests of the null that all pre-shock interactions equal zero. Semiannual event studies serve as the primary visualization because they reduce noise from small quarterly cells. I also implement placebo timing exercises that re-center the shock date at alternative pre-period points, and randomization inference that permutes block-level treatment assignments. These diagnostics, together with boundary designs, alternative fixed-effects structures, and sensitivity to the violence exposure definition, are reported in Section 8.

5 Results

5.1 Prices and Rents

Table 3 presents the main difference-in-differences estimates for residential sale prices (Panel A) and asking rents (Panel B). Columns (1) and (2) use the categorical treatment indicator, which distinguishes low exposure ($V_b = 1$) from severe exposure ($V_b \geq 2$) with unexposed blocks ($V_b = 0$) as the reference group. Columns (3) and (4) use the standardized continuous intensity $\tilde{V}_b^{std}$. Columns (2) and (4) augment the baseline with accessibility-quintile-by-year fixed effects, which allow flexible differential trends across levels of pre-shock GTFS centrality.

Because violence is spatially concentrated in central, well-connected neighborhoods, the accessibility quintiles are correlated with treatment intensity, so quintile-by-year fixed effects absorb part of the identifying variation alongside any differential centrality trends. Columns (2) and (4) should therefore be read as conservative. In practice the two specifications deliver similar answers, as the post-normalization residential estimate moves from -5.6% to -4.7% and remains significant at the 1% level, and the continuous estimate from -1.7% to -1.5%. Residual post-COVID repricing of central locations, running along the same spatial gradient as violence exposure, therefore does not appear to account for the result. I return to the centrality confound directly in the triple-difference analysis of Section 4.2.

Sale prices. The categorical specification in column (1) reveals two patterns. First, low-exposure blocks ($V_b = 1$) show no price response at any horizon, since coefficients are within half a percentage point of zero across all three post-shock phases. Second, severely exposed blocks ($V_b \geq 2$) decline by 3.0% during the metro closure period, remain about 2.4% below controls during the COVID transition, and fall to 5.6% in the post-normalization period. The effect is therefore not a transitory disruption that decays. It is smallest while the metro network was down and largest once metro service had resumed and pandemic restrictions had ended, three to six years after the shock. This timing is difficult to reconcile with a purely temporary disruption story and is examined directly in Section 4.2.

Adding accessibility controls in column (2) leaves the pattern intact, with a post-normalization estimate of -4.7%. The continuous specification tells the same story through a within-category gradient: a one-standard-deviation increase in violence intensity is associated with a 0.9% price decline during the metro closure, 0.7% during the COVID transition, and 1.7% in the post-normalization period, attenuating only to 1.5% once accessibility controls are added. Both the categorical threshold and the continuous gradient thus point to a discount that deepens over time rather than one that fades.

Asking rents. Rents behave differently, since severely exposed blocks show rent declines of 2.0% during the metro closure period and 3.5% during the COVID transition, but only 1.5% in the post-normalization period, and this last estimate is significant only at the 10% level and is not distinguishable from zero (-0.9%) once accessibility controls are added. The continuous specification shows a similar patttern, where the per-standard-deviation rent gradient peaks at 1.1% during the COVID transition and falls back to 0.7% afterwards. Low-exposure blocks show no rent decline at any horizon; their post-normalization coefficient is in fact positive and marginally significant in column (2), a pattern I do not press further since it appears only with accessibility controls.

In the semiannual event-time specification, the joint test that all pre-shock violence interactions equal zero yields $p = 0.260$ for rents, and in Section 8, both exposure levels pass comfortably the test. However, there is a caveat as these are asking rents posted on TOCTOC, the leading Chilean real estate listings portal, rather than transacted contract rents, so they need not track realized rents one for one if landlords adjust concessions or vacancy duration instead of headline rents.

Lastly, in the post-normalization period, severely exposed blocks trade at a discount of 5.6% while their asking rents stand only about 1.5% below controls, and the same asymmetry appears in the continuous specification (-1.7% against -0.7% per standard deviation). Sale prices fall and stay down; rents recover most of their initial decline. That divergence is not a loose end but a measurable object, and I turn to it below.

Commercial prices. Commercial sale prices are imprecisely estimated and statistically indistinguishable from zero in every specification and phase (Appendix Table A.7). Point estimates fluctuate in sign across phases—positive during the metro closure and the COVID transition, and essentially zero in the post-normalization period—with standard errors four to five times larger than those for residential prices. The imprecision reflects the smaller sample (29,295 transactions) and the sparseness of commercial trading at the block level, where only 24% of severely exposed blocks record any commercial transaction over the sample period, against 87% for residential; I return to it in Section 5.4, where I examine transaction activity as an alternative margin of adjustment.

Magnitudes. To gauge the scale of the residential effect, consider the post-normalization discount for severely exposed blocks. In the pre-shock period, the median residential sale price among blocks with $V_b \geq 2$ is roughly 2,500 UF, so a discount of 5.6% implies a capitalized loss of about 141 UF per property. Applied to the 298,648 residential units located in the 4,215 severely exposed blocks, this corresponds to an implied revaluation of approximately 42.1 million UF, or between USD 1.6 and 1.7 billion depending on the reference exchange rate. Using mean rather than median pre-shock prices raises the figure to 52.2 million UF, or roughly USD 2.0–2.1 billion; I report the median-based number as the conservative headline. Two qualifications are important. The calculation applies the estimated discount to the entire exposed residential stock whether or not a given unit transacted, so it measures a capitalized change in valuations rather than a sum of realized transaction losses. And it is a revaluation, not a welfare calculation: it does not net out the corresponding gain to buyers, nor does it account for the residents who left the exposed blocks entirely (Section 5.5).

Residential cap rates. A price decline that exceeds the contemporaneous decline in rents mechanically raises the implied gross yield on the exposed asset. Because sale and rental observations are not the same units, I estimate the wedge in a stacked specification that pools both outcomes and recovers the difference between the rent and price coefficients for the same exposure group and period, with standard errors accounting for the joint estimation.

The block-level and census-zone estimates differ in the coverage they can achieve. Only 31% of blocks record both a sale and a rental listing, against 85% of census zones, so I report the census-zone estimate as the lead and the block-level estimate as a robustness check. Full results by family, spatial unit, and period are in Table 4.

In the post-normalization period, severe exposure raises the implied capitalization rate by 4.8 log points. On a baseline gross yield of about 6.2%, this is equivalent to a rise from 6.2% to 6.5%, or roughly 0.3 percentage points. The estimate is stable across the three specification families used elsewhere in the paper—the violence-only design, the GTFS-based triple difference, and the binary metro triple difference—with point estimates between 4.2 and 5.4 log points and t-statistics between 2.5 and 3.0 in every case.

The interpretation requires care, because two distinct stories generate the same wedge. Firstly, we can talk about risk, as if severe exposure makes a location's future income stream more uncertain, investors require a higher return to hold the asset, and the price falls relative to current rents without any change in expected cash flows. The second is about fundamentals, since if market participants expect rents in exposed blocks to grow more slowly than elsewhere, prices capitalize that expectation immediately while current asking rents do not yet reflect it. These are not separately identified by the wedge alone, and I do not attempt to force a choice between them. The evidence elsewhere in the paper is, however, more consistent with the fundamentals channel being at least partly active. Commercial market participation falls persistently in the same blocks (Section 5.4), and exposed blocks lose population between 2017 and 2024 (Section 5.5), both of which point to a deterioration in the local demand conditions that support future rents rather than to pure repricing of a stable income stream. What the wedge does establish is that the residential adjustment is not a simple downward shift in the level of local housing services: prices and rents move by different amounts, and the difference is economically meaningful and statistically robust. The wedge is also more robust than either of its components. Adding CBD-bin-by-year fixed effects, the most demanding of the spatial controls considered in Section 8, attenuates the price effect to -4.2% and moves the rent effect to a precise zero, yet the implied difference between them is 4.4 log points, essentially unchanged from the baseline estimate of 4.8.

Figure 6 reports semiannual event-study estimates for residential sale prices, asking rents, and commercial sale prices, separately for $V_b = 1$ and $V_b \geq 2$. Pre-treatment coefficients are jointly insignificant in all three markets. The cleanest dynamic pattern appears in residential prices: low-

Table 3: Difference-in-Differences Estimates: Prices and Rents

	(1) Cat.	(2) Cat + Access	(3) Cont.	(4) Cont. + Access
Panel A: Residential Sale Prices (N=619,710)				
$V_b = 1$ × Metro closure	−0.006	0.000		
	(0.012)	(0.012)		
$V_b = 1$ × COVID transition	−0.002	0.004		
	(0.010)	(0.010)		
$V_b = 1$ × Post-normalization	−0.005	0.002		
	(0.013)	(0.014)		
$V_b \geq 2$ × Metro closure	−0.030**	−0.023*		
	(0.012)	(0.012)		
$V_b \geq 2$ × COVID transition	−0.024*	−0.015		
	(0.013)	(0.013)		
$V_b \geq 2$ × Post-normalization	−0.056***	−0.047***		
	(0.013)	(0.013)		
$\tilde{V}_b^{std}$ × Metro closure			−0.009***	−0.008***
			(0.003)	(0.003)
$\tilde{V}_b^{std}$ × COVID transition			−0.007**	−0.006*
			(0.003)	(0.003)
$\tilde{V}_b^{std}$ × Post-normalization			−0.017***	−0.015***
			(0.004)	(0.004)
Panel B: Residential Asking Rents (N=648,054)				
$V_b = 1$ × Metro closure	−0.006	0.000		
	(0.009)	(0.009)		
$V_b = 1$ × COVID transition	−0.004	0.003		
	(0.012)	(0.012)		
$V_b = 1$ × Post-normalization	0.016	0.022**		
	(0.010)	(0.010)		
$V_b \geq 2$ × Metro closure	−0.020**	−0.014*		
	(0.008)	(0.008)		
$V_b \geq 2$ × COVID transition	−0.035***	−0.028***		
	(0.008)	(0.008)		
$V_b \geq 2$ × Post-normalization	−0.015*	−0.009		
	(0.009)	(0.009)		
$\tilde{V}_b^{std}$ × Metro closure			−0.007***	−0.006***
			(0.002)	(0.002)
$\tilde{V}_b^{std}$ × COVID transition			−0.011***	−0.010***
			(0.002)	(0.002)
$\tilde{V}_b^{std}$ × Post-normalization			−0.007***	−0.006***
			(0.002)	(0.002)
Block FE	Yes	Yes	Yes	Yes
Year-month FE	Yes	Yes	Yes	Yes
Municipality × year FE	Yes	Yes	Yes	Yes
Accessibility quintile × year FE	No	Yes	No	Yes
R^2 (Panel A)	0.844	0.844	0.844	0.844
R^2 (Panel B)	0.627	0.627	0.627	0.627

Notes: Dependent variable: ln(price in UF) for sale prices (Panel A) and ln(UF/m^2) for asking rents (Panel B). All specifications include property controls (log construction area, log lot size, log year built, property type). Metro closure: Oct 2019–Sep 2020. COVID transition: Oct 2020–Sep 2022. Post-normalization: Oct 2022 onward. The categorical specification (Columns 1–2) uses $V_b = 0$ as the reference group. The continuous specification (Columns 3–4) uses the standardized continuous intensity $\tilde{V}_b^{std}$ (mean zero, unit variance among $V_b > 0$). Accessibility quintiles are based on pre-shock GTFS travel time to the CBD. Standard errors clustered at the block level in parentheses.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 4: Implied Capitalization Rates: Sale Prices, Asking Rents, and the Wedge

	(1) Sale price	(2) Asking rent	(3) Implied cap rate (rent − sale)	(4)	(5)
	Coef.	Coef.	Coef.	SE	t
Panel A. Census zone (lead specification)					
Low exposure ($V_b = 1$)					
Metro closure	0.001	−0.003	−0.004	0.021	−0.17
	(0.020)	(0.011)			
COVID transition	−0.008	−0.005	0.003	0.025	0.11
	(0.024)	(0.012)			
Post-normalization	0.005	0.022**	0.018	0.018	1.01
	(0.017)	(0.011)			
Severe exposure ($V_b \geq 2$)					
Metro closure	−0.004	0.004	0.008	0.018	0.46
	(0.017)	(0.009)			
COVID transition	−0.003	−0.017*	−0.014	0.018	−0.80
	(0.017)	(0.009)			
Post-normalization	−0.036**	0.012	**0.048**	0.018	**2.71**
	(0.016)	(0.011)			
Panel B. Census block (robustness)					
Low exposure ($V_b = 1$)					
Metro closure	−0.008	−0.008	0.001	0.015	0.05
	(0.013)	(0.008)			
COVID transition	−0.005	−0.010	−0.005	0.015	−0.33
	(0.011)	(0.011)			
Post-normalization	−0.008	0.016	0.024	0.016	1.49
	(0.013)	(0.011)			
Severe exposure ($V_b \geq 2$)					
Metro closure	−0.025**	−0.012	0.014	0.015	0.93
	(0.012)	(0.008)			
COVID transition	−0.023	−0.035***	−0.012	0.016	−0.77
	(0.014)	(0.008)			
Post-normalization	−0.055***	−0.007	**0.048**	0.016	**2.96**
	(0.014)	(0.010)			
Panel C. Post-normalization wedge for $V_b \geq 2$, across specification families					
			Coef.	SE	t
Violence only, census zone	−0.036	0.012	0.048	0.018	2.71
Violence only, block	−0.055	−0.007	0.048	0.016	2.96
GTFS triple difference, census zone	−0.048	−0.004	0.045	0.018	2.47
GTFS triple difference, block	−0.058	−0.016	0.042	0.016	2.54
Metro triple difference, census zone	−0.041	0.012	0.054	0.020	2.73
Metro triple difference, block	−0.048	−0.005	0.043	0.017	2.51
Unit × outcome FE			Yes		
Year-month FE			Yes		
Municipality × year FE			Yes		

Notes. Sale and rental observations are pooled into a single stacked dataset in which each observation carries an indicator for its outcome type, and the full specification is interacted with that indicator. Columns (1) and (2) therefore report the sale-price and asking-rent responses estimated jointly on the same sample and fixed-effects structure, and column (3) reports the difference between them, which is the effect on the implied gross capitalization rate. Because the two coefficients come from one estimation, the standard error in column (4) accounts for their covariance rather than treating them as independent. All quantities are in log points; a value of 0.048 corresponds to a 4.8% proportional increase in the implied yield, which on a baseline gross yield of 6.2% amounts to approximately 0.31 percentage points. Standard errors clustered at the spatial unit. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

exposure blocks remain flat throughout the sample, while severely exposed blocks decline gradually and without reversal, with the largest coefficients in the final semesters. Asking rents for $V_b \geq 2$ move in the same direction initially, reach their trough during the COVID transition, and then revert toward zero, so the price–rent divergence documented above is visible in the dynamics and not only in the phase averages. Commercial estimates are noisy throughout and do not display a clear post-shock trajectory at either exposure level; their pre-shock test passes, but with confidence intervals several times wider than in the residential panel this offers limited reassurance.

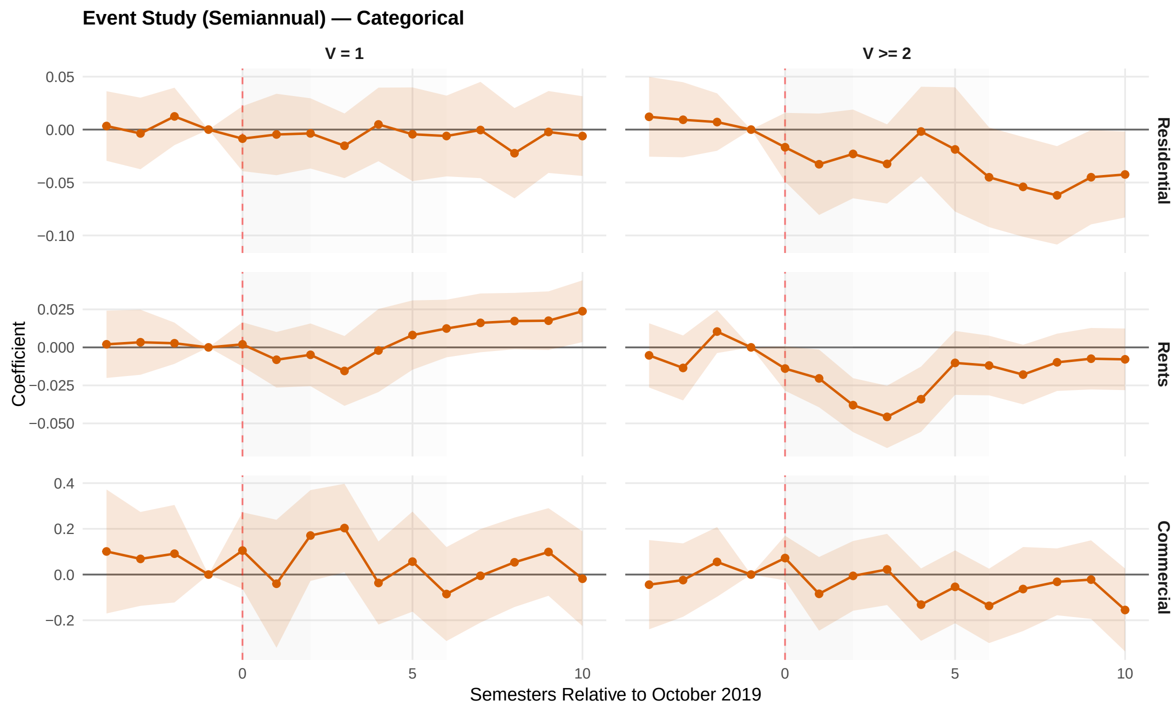

Figure 6: Event Study (Semiannual): Prices and Rents by Violence Exposure. Coefficients from semiannual event-study specifications interacting half-year indicators with $V_b = 1$ (left) and $V_b \geq 2$ (right), with $V_b = 0$ as the reference group. Rows: residential sale prices (top), asking rents (middle), commercial sale prices (bottom). Reference period: semester -1. Shaded areas denote the metro closure (light) and COVID transition (medium) phases. All models include block, year-month, and municipality-by-year fixed effects. 95% CIs based on SEs clustered at the block level. Joint Wald tests of the null that all pre-shock interactions equal zero: $p = 0.958$ (residential sale prices), $p = 0.260$ (asking rents), $p = 0.649$ (commercial sale prices).

5.2 Spillovers and boundary comparisons

Two opposite concerns arise from the way the control group is defined. The first is contamination, as unexposed blocks adjacent to hotspots may not be truly unexposed, since stigma, foot-traffic reallocation, and visible disorder can extend beyond the prosecutorial polygons used to construct V_b. If nearby $V_b = 0$ blocks are partially affected, the baseline compares severely exposed blocks against a contaminated control group and attenuates the estimate toward zero. The second is comparability, as if the effective comparison is between exposed blocks and distant, structurally different parts of the city, the estimate may reflect broad spatial heterogeneity rather than local exposure. Table 5 addresses both, first by locating the effect in space and then through two symmetric exercises around the hotspot boundary.

Panel A decomposes the post-normalization effect by distance to the nearest severely exposed block, taking controls more than 1,500m away as the reference. Severely exposed blocks themselves show a discount of 5.7%, essentially identical to the baseline estimate. No ring of unexposed blocks differs significantly from distant controls, including the two nearest ones. The point estimates in those rings are negative but small relative to the treated effect and imprecisely estimated (−1.5% within 250m and −2.4% between 250 and 500m), so modest short-range spillover cannot be ruled out; if present, it would bias the baseline toward zero rather than away from it. On price, therefore, the stable-unit assumption is not obviously violated and the control group appears clean. Panels B and C confirm this from both directions. Dropping observations near the boundary enlarges the discount, from −5.6% in the full sample to −6.7% once blocks within 250m of the edge are excluded, which is what contamination of proximate controls would produce. Restricting the sample to a narrow band around the boundary instead preserves it, at between −4.4% and −5.4% across bandwidths from 1,000 to 2,500 metres, always significant at the 1% level. Low-exposure blocks remain indistinguishable from controls in every bandwidth.

The estimates in the rest of the paper should therefore be read as conservative. One qualification is worth stating, since this evidence disciplines the price estimates, not the reach of the shock itself. Population responses do extend into the immediate ring around exposed blocks even where prices there do not move, a contrast I return to in Section 5.5.

5.3 Violence versus Accessibility

The baseline estimates in Table 3 cannot by themselves distinguish a violence-driven revaluation of place from metro-accessibility disruptions, because the two shocks are spatially correlated. This subsection separates them using the channel decomposition and triple-difference specifications introduced in Section 4.2.

Table 5: Spillovers and boundary designs: residential sale prices, post-normalization

	(1)	(2)	(3)	(4)	(5)
Panel A. Distance rings (reference: unexposed controls beyond 1,500m)					
	1,000–1,500m	500–1,000m	250–500m	0–250m	$V_b \geq 2$
Post-normalization	0.015	−0.013	−0.024	−0.015	-0.057^{***}
	(0.012)	(0.017)	(0.021)	(0.026)	(0.016)
Block-transactions	79,903	80,143	39,628	24,247	102,855
Panel B. Donut: drop observations within a distance of the edge					
Drop within	None	100m	250m	500m	
$V_b \geq 2\times$ Post-normalization	-0.056^{***}	-0.056^{***}	-0.067^{***}	-0.067^{***}	
	(0.013)	(0.013)	(0.014)	(0.016)	
Observations	619,710	589,058	542,569	461,172	
Panel C. Bandwidth: keep observations within a distance of the edge					
Keep within	1,000m	1,500m	2,000m	2,500m	
$V_b \geq 2\times$ Post-normalization	-0.044^{***}	-0.052^{***}	-0.051^{***}	-0.054^{***}	
	(0.015)	(0.013)	(0.013)	(0.013)	
$V_b = 1\times$ Post-normalization	−0.003	−0.003	−0.003	−0.003	
	(0.015)	(0.014)	(0.013)	(0.013)	
Observations	305,242	408,499	479,211	517,399	
Block FE	Yes	Yes	Yes	Yes	
Year-month FE	Yes	Yes	Yes	Yes	
Municipality × year FE	Yes	Yes	Yes	Yes	

Notes. The dependent variable is ln(price in UF) for residential sale transactions. Panel A assigns each unexposed block to a ring based on the distance from its centroid to the nearest block with $V_b \geq 2$, and reports post-normalization coefficients relative to unexposed blocks beyond 1,500m; the last column reports the effect for severely exposed blocks themselves. Panels B and C use the signed distance from each block to the nearest boundary between exposed and unexposed areas: Panel B drops observations within the indicated distance of that boundary, with column (1) as the full-sample baseline, and Panel C retains only observations within it. All specifications include property controls and the baseline fixed effects. Standard errors clustered at the block level.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Channel decomposition. Table 6 reports estimates from the six-group channel specification, which crosses violence levels ($V_b = 0$, $V_b = 1$, $V_b \geq 2$) with metro proximity ($M_b = 0$, $M_b = 1$). The reference category is unexposed blocks far from closed stations. I focus on the post-normalization period, when metro service has resumed and temporary mobility disruptions should no longer dominate the results.

Three patterns emerge. First, *Metro Only* blocks ($V_b = 0$, near a closed station) show no persistent decline in either market. The post-normalization coefficient is +3.1% for sale prices, marginally significant with the wrong sign for an accessibility story ($p = 0.096$), and +0.6% for rents ($p = 0.772$). These blocks lost station access for the same eleven months as the exposed ones but recorded no violence, and five years later they trade at no discount whatsoever. Temporary transit disruption on its own is therefore not sufficient to generate persistent residential price effects.

Second, *Violence Only* blocks ($V_b \geq 2$, far from closed stations) show a large and precisely estimated decline in sale prices: −5.0% in the post-normalization period. Because these blocks were not directly exposed to station closures, their persistent discount is difficult to attribute to accessibility and points to violence itself as the driver of the residential response. Rents behave differently, as in the baseline, since the Violence Only coefficient is −0.4% and indistinguishable from zero, even though the same cell shows a clear decline of 2.9% during the COVID transition. The channel decomposition therefore reproduces the price–rent divergence of Section 5.1 within the cell where violence operates in isolation, which is the cleanest place to observe it.

Third, blocks exposed to *Both* severe violence and metro proximity display more negative point estimates for sale prices than Violence Only blocks (−7.3% versus −5.0%), but the incremental difference of 2.3 percentage points is not statistically significant; the same holds for rents, where the increment is 0.6 percentage points. The point estimates are compatible with some compounding across channels, but they do not provide robust evidence that metro proximity systematically amplifies the effect of severe violence, and I do not treat compounding as an identified result.

Low-exposure blocks ($V_b = 1$) show no persistent sale-price effect: −0.3% far from closed stations and essentially zero near them. The severity gradient is nonetheless sharp. The difference between severely and low-exposed blocks far from closed stations is −4.7 percentage points for sale prices and −2.3 percentage points for rents, reinforcing the threshold pattern documented in Section 5.1. In rents the gradient reflects a positive coefficient for low exposure (+2.0%) as much as a negative one for severe exposure, so it should be read as a relative ranking rather than as evidence of a persistent rent discount.[11]

Triple-difference. Table 7 re-parameterizes the comparison using two continuous measures of transit disruption. Column (1) uses the GTFS-based network accessibility shock—the increase in travel time to the CBD when closed stations are removed from the network—and column (2) uses a Euclidean measure of station-access disruption. Both are scaled by their standard deviation without centering, so that zero corresponds to no disruption. This scaling is what makes the specification informative: the main effect of $\mathbb{1}\{V_b \geq 2\}$ is then the violence discount evaluated at zero accessibility loss, that is, a pure Violence Only estimate recovered from the full sample rather than from a single cell. I do not report a binary-M_b triple difference in the main text, since it is an exact re-parameterization of the channel decomposition already shown in Table 6.

Both measures deliver the same three results. The violence main effect is negative, large, and precisely estimated for sale prices of about -6.3% under either measure. Accessibility disruption on its own does not depress prices; the GTFS coefficient is in fact positive and marginally significant (+1.4%), and the station-distance coefficient is a precise zero. And the interaction is *positive* and significant—+2.5 percentage points per standard deviation of network shock, +0.7 percentage points under station distance—meaning that among severely exposed blocks, those that lost more accessibility display smaller discounts. This is the opposite of what an accessibility-driven account predicts, and it is not confined to the post-normalization period: the interaction is also positive and significant during the metro closure itself (+2.4 percentage points). Rents follow the same signs with muted magnitudes, as a violence main effect of -1.6%, no accessibility effect, and an interaction indistinguishable from zero.

Why the interaction is positive is not something the data settle. One possibility is that the blocks with the largest network shocks are those whose access depended most heavily on a single station, and that their recovery once service resumed was correspondingly complete. What matters for the argument is the sign: no weighting of a pure accessibility channel reproduces a pattern in which disruption is associated with a milder discount.

[11] Commercial price channel estimates are reported in Appendix Table A.9. All post-normalization coefficients are imprecise and show no clear pattern; the one significant estimate anywhere in the commercial channel table is a positive coefficient for severely exposed blocks near closed stations during the metro closure period itself (+14.5%), which does not survive into later phases.

Table 6: Channel decomposition estimates for prices and rents

	(1) Sale Prices	(2) Asking Rents
Metro closure period		
$V_b = 0$, near metro	-0.002	0.023
	(0.019)	(0.022)
$V_b = 1$, far	-0.005	-0.008
	(0.012)	(0.009)
$V_b = 1$, near metro	-0.017	0.020
	(0.040)	(0.030)
$V_b \geq 2$, far	-0.032***	-0.015*
	(0.012)	(0.008)
$V_b \geq 2$, near metro	-0.021	0.012
	(0.021)	(0.016)
COVID transition		
$V_b = 0$, near metro	0.007	0.002
	(0.017)	(0.020)
$V_b = 1$, far	0.001	-0.009
	(0.011)	(0.011)
$V_b = 1$, near metro	-0.028	0.003
	(0.025)	(0.020)
$V_b \geq 2$, far	-0.017	-0.029***
	(0.014)	(0.009)
$V_b \geq 2$, near metro	-0.055***	-0.046***
	(0.020)	(0.012)
Post-normalization		
$V_b = 0$, near metro	0.031*	0.006
	(0.019)	(0.020)
$V_b = 1$, far	-0.003	0.020*
	(0.014)	(0.011)
$V_b = 1$, near metro	-0.000	0.000
	(0.025)	(0.018)
$V_b \geq 2$, far	-0.050***	-0.004
	(0.013)	(0.011)
$V_b \geq 2$, near metro	-0.073***	-0.009
	(0.020)	(0.013)
Block FE	Yes	Yes
Year-month FE	Yes	Yes
Municipality × year FE	Yes	Yes
Observations	619,710	648,054
R^2	0.844	0.412

Notes: Six-group channel decomposition crossing violence exposure ($V_b = 0$, $V_b = 1$, $V_b \geq 2$) with metro proximity ($M_b = \mathbb{1}\{\text{dist. to closed station} \leq 500\text{m}\}$). Reference group is $V_b = 0$, far from closed stations. All specifications include property controls. Formal post-normalization contrasts are discussed in the text. Standard errors clustered at the block level.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

This specification also carries the identification weight in the paper. Its pre-treatment interaction terms are jointly insignificant for both residential prices and rents, whereas the corresponding test on the baseline violence terms rejects at monthly frequency for severe exposure (Section 8). I therefore treat the triple difference, rather than the two-way baseline, as the specification on which the violence-versus-accessibility conclusion rests. The commercial counterpart fails its pre-trend test decisively and is reported in the Appendix without an identification claim. Full results across all periods, for both continuous measures and for commercial properties, are in Appendix Tables A.11–A.12.

Table 7: Triple-difference estimates in the post-normalization period

	(1) GTFS network shock	(2) Station distance
Panel A. Residential sale prices		
$\mathbb{1}\{V_b \geq 2\}\times$ Post-normalization	-0.063^{***}	-0.063^{***}
	(0.015)	(0.016)
Accessibility shock $\times$ Post-normalization	0.014^{*}	0.003
	(0.008)	(0.004)
$\mathbb{1}\{V_b \geq 2\}\times$ Accessibility shock $\times$ Post-normalization	0.025^{**}	0.007^{**}
	(0.010)	(0.003)
Observations	619,710	619,710
R^2	0.844	0.844
Panel B. Residential asking rents		
$\mathbb{1}\{V_b \geq 2\}\times$ Post-normalization	-0.016^{*}	-0.017^{**}
	(0.008)	(0.008)
Accessibility shock $\times$ Post-normalization	-0.003	-0.002
	(0.019)	(0.007)
$\mathbb{1}\{V_b \geq 2\}\times$ Accessibility shock $\times$ Post-normalization	0.006	0.006
	(0.018)	(0.008)
Observations	648,054	648,054
R^2	0.411	0.411
Block FE	Yes	Yes
Year-month FE	Yes	Yes
Municipality $\times$ year FE	Yes	Yes
Pre-trend Wald p (triple, prices)	0.624	
Pre-trend Wald p (triple, rents)	0.478	

Notes. The table reports coefficients for interactions with the post-normalization period; coefficients for earlier periods are estimated but not shown (see Appendix Tables A.11 and A.12 for complete results). Column (1) uses the GTFS-based network accessibility shock, defined as the increase in shortest-path travel time to the CBD when closed stations are removed from the transit network. Column (2) uses a Euclidean measure of station-access disruption, correlated but not identical to the GTFS network shock ($\rho = 0.57$). Both measures are scaled by their standard deviation without centering, so that zero corresponds to no network disruption and the main effect of $\mathbb{1}\{V_b \geq 2\}$ captures the violence discount for blocks with no accessibility loss. A binary-proximity version of this specification is an exact re-parameterization of the channel decomposition in Table 6 and is not reported separately. Pre-trend p-values are joint Wald tests on the pre-shock triple interaction terms in the semiannual event-time specification. All specifications include property-level controls. Standard errors clustered at the block level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

5.4 Transaction Activity: Rotation and Market Thinning

Table 8 reports the main estimates for transaction activity. I focus on the extensive margin, defined as whether a block records any transaction in a given quarter, because this is the clearest measure of market participation and the least sensitive to selection among the subset of blocks that remain active. Appendix Tables A.13 and A.14 report the corresponding count and turnover specifications.

Residential activity. The residential market does not thin. Severely exposed blocks ($V_b \geq 2$) are no less likely to transact during the metro closure or the COVID transition, and in the post-normalization period they are more likely to do so, by 0.55 percentage points. Against a pre-shock control-block baseline of 12.2%, this is a relative increase of about 4.5%, or roughly 273 additional transacting block-quarters. Read alongside Table 3, the combination informs that prices fall by 5.6% while transactions become more frequent. A market that stops functioning produces the opposite pairing, with prices held up by the absence of trade. What the residential data show instead is consistent with rotation—owners exiting exposed blocks at lower prices, and buyers willing to transact at those prices—so that the discount is realized rather than merely posted. The positive coefficient is the least robust element of the table: it attenuates to 0.22 percentage points and loses significance once accessibility-quintile-by-year fixed effects are added, and to 0.45 percentage points with CBD-distance controls (Appendix Table A.13). I therefore treat the direction as suggestive and the absence of contraction as the firm result.

Commercial activity. The commercial market behaves differently. Blocks with severe violence exposure are significantly less likely to record any transaction in all three post-shock phases. The effect is largest during the metro closure period (-1.64 percentage points), remains negative during the COVID transition (-0.73 percentage points), and persists into the post-normalization period (-0.78 percentage points). The continuous specification is negative and significant throughout. Unlike its residential counterpart, this estimate is unaffected by centrality controls, remaining at -0.74 and -0.69 percentage points and significant at the 1% level under accessibility-quintile and CBD-distance controls respectively. This is the clearest quantity response in the paper. Violence does not merely lower the value of commercial locations; it reduces the probability that they transact at all.

Counts and selection among surviving blocks. The count specifications in the appendix refine this picture. In residential markets, PPML estimates for transaction counts are positive but imprecise, pointing in the same direction as the extensive margin without adding statistical content. In commercial markets, PPML counts do not fall in parallel with the extensive margin and are weakly positive in the post-normalization period. I interpret this not as evidence against thinning, but as evidence of concentration among the subset of blocks that continue to transact; i.e., violence exposure reduces participation at the block-quarter level, while activity among surviving active blocks need not decline proportionally. Turnover rates are small in both markets and do not alter this reading, though the residential turnover coefficient is marginally positive in the post-normalization period, again consistent with rotation rather than freezing.

Accessibility and transaction activity. Appendix Table A.15 reports triple-difference specifications for volume outcomes using closure-induced access disruption. These estimates do not support a purely accessibility-driven account of the extensive-margin results. In residential markets, all three terms—violence, accessibility loss, and their interaction—are statistically indistinguishable from zero, so participation is insensitive to both shocks. In commercial markets, the violence main effect remains negative and significant at zero disruption (-0.97 percentage points) while the interaction with the network shock is positive ($+0.49$ percentage points), reproducing on quantities the sign pattern documented for residential prices in Section 5.3. One qualification is worth reporting: under the Euclidean station-distance measure, accessibility loss does have an independent negative effect on commercial participation (-0.22 percentage points), so transit disruption is not irrelevant to commercial trading. Even there, however, the violence effect is larger and the interaction still positive, so the contraction cannot be attributed to lost accessibility alone.

The results point to different adjustment margins across markets. In residential real estate, violence is reflected in prices, since sale prices fall persistently while transaction activity, if anything, increases. The market remains active enough for the shock to be capitalized, and the observed price is a meaningful signal of the revaluation. In commercial real estate, the pattern inverts. Prices are noisy and statistically uninformative, while participation contracts sharply and persistently. The asymmetry is anticipated by the raw coverage of the two markets noted in Section 5.1: 87% of severely exposed blocks record a residential transaction over the sample period, against only 24%

for commercial. Where trading is frequent the shock appears in prices; where it is sparse it appears in whether trade occurs at all. The implication for measurement is direct: hedonic estimates based only on observed commercial transactions will understate the economic loss, because the blocks most affected are disproportionately the ones that drop out of the sample.

The commercial extensive-margin result is also economically meaningful. In the pre-shock period, the baseline probability that a control block-quarter records any commercial transaction is 2.18%. The post-normalization estimate of -0.78 percentage points therefore corresponds to a reduction of roughly 36% in the probability of any commercial transaction. Applied to the 1,810 severely exposed blocks that hold commercial stock over the twelve post-normalization quarters, this implies on the order of 169 fewer transacting block-quarters. This is best interpreted as a reduction in market participation rather than a direct count of missing transactions, but it fixes the scale of the commercial thinning documented in Table 8.

5.5 Displacement and Recomposition

Prices and transaction activity describe how the market responded; they do not say who ended up living in the exposed blocks. This matters for two reasons. If severely exposed blocks lost residents, the price discount is accompanied by a real reallocation of population rather than being a purely financial revaluation. And if the residents who left were systematically different from those who arrived, part of the measured discount could reflect a change in the composition of the neighbourhood rather than a change in the value of the location itself. I address both using block-level census data for 2017 and 2024, interpolated onto the 2017 block grid.[12] This evidence is exploratory as two cross-sections cannot support a pre-trend test, and the 2017–2024 window spans the unrest, the pandemic, and a large migration episode.

Depopulation. Panel A of Table 9 reports the change in total population between 2017 and 2024 by exposure. Severely exposed blocks lose population relative to unexposed blocks in the same comuna, by about 3%. The estimate is -2.9% in the log-ratio specification on stable blocks

[12]The 2024 census geography does not coincide with the 2017 one. I interpolate 2024 counts onto 2017 blocks and flag as stable those blocks with a dominant one-to-one correspondence between the two grids; rural comunas are excluded. Specifications with and without the stability restriction are reported separately. All models include comuna fixed effects, so identification comes from within-comuna variation in exposure, and standard errors are clustered at the block level.

Table 8: Transaction activity on the extensive margin

	Residential		**Commercial**	
	Categorical	Continuous	Categorical	Continuous
Metro closure period				
$V_b = 1$	−0.0001		−0.0062**	
	(0.0025)		(0.0025)	
$V_b \geq 2$	−0.0036		−0.0164***	
	(0.0031)		(0.0034)	
$\tilde{V}_b^{std}$		−0.0042***		−0.0068***
		(0.0013)		(0.0014)
COVID transition				
$V_b = 1$	0.0018		−0.0018	
	(0.0023)		(0.0025)	
$V_b \geq 2$	−0.0009		−0.0073**	
	(0.0027)		(0.0030)	
$\tilde{V}_b^{std}$		0.0005		−0.0032***
		(0.0012)		(0.0013)
Post-normalization				
$V_b = 1$	0.0040*		−0.0020	
	(0.0022)		(0.0022)	
$V_b \geq 2$	0.0055**		−0.0078***	
	(0.0027)		(0.0027)	
$\tilde{V}_b^{std}$		0.0021		−0.0034***
		(0.0013)		(0.0011)
Block FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Municipality × year FE	Yes	Yes	Yes	Yes
Observations	1,704,456	1,704,456	486,540	486,540
R^2	0.307	0.307	0.282	0.282

Notes. Dependent variable is an indicator for whether block b records any transaction in quarter t. Categorical specifications use $V_b = 0$ as the reference group. Continuous specifications use the standardized violence intensity $\tilde{V}_b^{std}$. All models include block fixed effects, quarter fixed effects, and municipality-by-year fixed effects. Standard errors clustered at the block level.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

and −3.0% when all blocks are included; the Poisson specification with a 2017 population offset gives −4.4% on the full sample and a weaker −2.2% on stable blocks. Three of four specifications are significant at the 1% level and all four are negative, so I read the result as robust in sign and around 3% in magnitude. Low-exposure blocks show no population change in any specification, reproducing the same threshold that appears in prices, since exposure to a single incident type does not move outcomes, exposure to multiple types does.

The magnitudes line up with the market evidence rather than substituting for it. A 5.6% price discount alongside a 3% population loss and a rise in transaction frequency (Section 5.4) describes blocks that people left and that cleared at lower prices, not blocks that became illiquid. The population response also reaches slightly further in space than the price response, since depopulation extends into the ring of unexposed blocks within 250m of a hotspot, whereas prices in those same rings are statistically indistinguishable from distant controls (Table 5, Panel A). Residential mobility spills across the boundary; capitalization does not. This is a useful asymmetry, because it means the control group can be contaminated on population while remaining clean for the purpose of estimating the price effect.

Composition. Panel B turns to who remained. Using the share of residents by educational attainment among those with known education, severely exposed blocks show a decline of 0.65 percentage points in the tertiary-educated share and a rise of 0.57 percentage points in the primary-educated share between 2017 and 2024 (columns 1 and 2). Taken at face value this looks like downward socioeconomic sorting. Two cautions apply before reading it that way. The measure covers the total resident population rather than household heads, so it captures who lives in the block and not the socioeconomic position of the households occupying it. And the outcome is a compositional share, so it reflects the mix of departures and arrivals, not any change in the attainment of individuals.

The average conceals opposite patterns, which is the more interesting result. Columns (3) and (4) interact exposure with the comuna's relative intensity of migratory inflow.The interaction is large and negative for both outcomes: −4.6 percentage points for the primary share and −4.7 for the tertiary share. The level terms in these columns are therefore not comparable to the pooled estimates, since they measure the effect at zero migratory intensity rather than at its average, which is why the tertiary coefficient switches sign between columns (2) and (4). Evaluated at the block-weighted mean migratory rank of 0.29, the interacted specification returns the pooled estimates exactly. Evaluated across the distribution, the implied effect of severe exposure on the primary-educated share moves from +1.4 percentage points in comunas at the tenth percentile of migratory intensity to −2.2 percentage points at the ninetieth. In low-inflow comunas, exposed blocks became less educated, which is the classic sorting pattern. In high-inflow comunas they became more educated. In low-inflow comunas, exposed blocks became less educated, which is the classic sorting pattern. The reading I find most plausible is that cheap, centrally located exposed blocks absorbed arrivals who were relatively well educated but had low incomes, which decouples educational attainment from the socioeconomic signal it usually carries in Chilean census data. I do not treat this as established, and it is the reason I do not report the pooled composition estimate as evidence of socioeconomic downgrading.

6 Mechanisms and Margins of Adjustment

The results in Section 5 point to a common underlying shock with different observable margins across markets. In residential real estate, violence is reflected in lower sale prices, in a rising implied yield, and in a net loss of residents, while transaction activity does not contract and if anything increases. This pattern is most consistent with a persistent revaluation of place rather than with a temporary disruption to accessibility. In commercial real estate, by contrast, prices are noisy and statistically uninformative, but participation falls sharply and persistently. The question, then, is not whether violence mattered in both markets, but why it is revealed through prices in one case and through market participation in the other.

I interpret this difference through the lens of market thickness. When trading is sufficiently active, valuation shocks are more likely to be capitalized into observed prices. When trading is sparse, the same shock may instead appear through non-trading, delayed transactions, or contraction in operating scale among incumbent firms. This section develops that interpretation in two steps. I first show that price adjustment is concentrated in thicker markets, which helps explain why the residential response is visible in prices while the commercial one is not. I then turn to firm-level administrative data and show that, in commercial areas, the adjustment operates through

Table 9: Population change and educational recomposition, 2017–2024

	(1)	(2)	(3)	(4)
Panel A. Total population				
	PPML Stable blocks	PPML All blocks	Log-ratio Stable blocks	Log-ratio All blocks
$V_b = 1$	0.003	−0.011	−0.000	−0.003
$V_b \geq 2$	−0.022*	−0.044***	−0.030***	−0.030***
	(0.012)	(0.013)	(0.007)	(0.008)
Observations	40,399	42,203	40,399	42,203
Panel B. Educational composition				
	Pooled		× Migratory rank	
	Δ Primary	Δ Tertiary	Δ Primary	Δ Tertiary
$V_b = 1$	—	—	0.0166***	−0.0026
			(0.0021)	(0.0021)
$V_b \geq 2$	0.0057***	−0.0065***	0.0190***	0.0073***
	(0.0016)	(0.0016)	(0.0022)	(0.0022)
$V_b = 1\times$ Migratory rank			−0.0285***	−0.0058
			(0.0052)	(0.0055)
$V_b \geq 2\times$ Migratory rank			−0.0459***	−0.0471***
			(0.0052)	(0.0054)
Observations	40,399	40,399	39,784	39,784
Comuna FE	Yes	Yes	Yes	Yes

Notes. Census blocks in urban comunas of the Santiago Metropolitan Region, with 2024 counts areally interpolated onto the 2017 block grid. *Stable* blocks are those with a dominant one-to-one correspondence between the two grids. Panel A: columns (1)–(2) report Poisson pseudo-maximum-likelihood estimates of 2024 population with $\log(\text{population}_{2017})$ as offset; columns (3)–(4) report OLS estimates with $\log(\text{pop}_{2024}+1)-\log(\text{pop}_{2017}+1)$ as the dependent variable. Standard errors for $V_b = 1$ are omitted for space and are of the same order as those reported for $V_b \geq 2$. Panel B: the dependent variable is the 2017–2024 change in the share of residents with primary or tertiary attainment, among residents with known attainment; the sample is restricted to stable blocks with at least twenty such residents in both years. Migratory rank is the comuna's zero-to-one rank of accumulated residence permits per 2017 resident; its main effect is absorbed by the comuna fixed effects. Pooled composition estimates without the interaction are +0.0057 (0.0016) for the primary share and −0.0065 (0.0016) for the tertiary share. Standard errors clustered at the block level.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

the thinning of establishments with employees rather than through firm exit.

6.1 Market thickness and the margin of adjustment

The contrast between residential and commercial outcomes suggests that the visibility of the shock depends on pre-shock market thickness. In markets with frequent transactions, shocks are more likely to appear in observed prices because trades continue to occur and price discovery remains active. In thinner markets, observed prices become less informative because the main response may occur at the extensive margin, through fewer transactions rather than lower prices conditional on trade.

Testing this requires comparing price and quantity margins across thickness levels within the same markets. I classify residential blocks as thick or thin according to whether their pre-shock transaction count, measured between the start of the sample and September 2019, lies above or below the median among blocks with at least one sale, and estimate the post-normalization effect separately in each cell for both prices and the extensive margin. Panel A of Table 10 reports the resulting two-by-two.

The residential pattern is unambiguous on the price margin. In thick blocks, severe violence lowers sale prices by 5.7%. In thin blocks, the corresponding estimate is -0.3% and indistinguishable from zero. The difference between the two, 5.4 percentage points, is significant at the 10% level, so thickness predicts where repricing occurs within a single market and not merely across markets. The quantity margin behaves differently: neither cell shows a significant extensive-margin response and the two are statistically indistinguishable from each other. Within residential real estate, therefore, thickness governs the price margin and leaves the quantity margin flat, which is precisely the asymmetry the mechanism predicts.

The commercial market cannot be split the same way. The median pre-shock commercial transaction rate across blocks is zero, as commercial trading is not merely thinner than residential trading, it is uniformly thin, so there is no interior variation in thickness to condition on. Forcing the split produces a degenerate comparison. The thin price cell contains 4,013 observations against 26,883 in the thick cell, and its coefficient is estimated with a standard error of over 30,000. I therefore report the commercial price effect pooled, at $+0.8\%$ with a standard error of 4.6 percentage points, and the commercial extensive-margin effect pooled at -0.78 percentage points ($p < 0.01$)

from Section 5.4. The two-by-two thus completes itself asymmetrically: residential real estate has both cells and the price effect lives in the thick one; commercial real estate has only the thin cell, and its adjustment runs entirely through participation.[13]

This interpretation has an important implication for measurement. In thicker markets, hedonic price regressions recover a significant share of the loss because transactions continue and prices adjust. In thinner markets, conditioning on observed trades understates the economic cost of violence, because the shock reveals itself partly through non-trading. The commercial results should therefore be read as evidence that price-based estimates alone are incomplete, not as evidence that commercial values were unaffected. The raw coverage figures make the point concrete: 87% of severely exposed blocks record a residential transaction over the sample period, against 24% for commercial. Where the market is thick enough to reprice, it reprices; where it is not, the adjustment shows up in whether trade happens at all.

6.2 Firm-level evidence on commercial retrenchment [Still in process]

The commercial transaction results show that violence reduces market participation rather than generating clean price adjustments. A natural follow-up question is whether this pattern reflects broader changes in local commercial activity. If violence reduces pedestrian demand and worsens neighborhood conditions for street-level businesses, exposed areas should show symptoms of commercial retrenchment beyond the real estate market itself.

To explore this margin I use administrative records from Chile's *Servicio de Impuestos Internos* (SII), which link registered firms to their reported workforce.[14] I build a balanced block–year panel from the cohort of firms geocoded before October 2019, covering 2015–2024, with 2018 as the reference year and 2019 treated as a transition year. Firms are classified by registered activity into four groups fixed before the shock: *direct local consumption* (specialized retail, restaurants and bars, personal services), *upstream local* (wholesale and input suppliers to those activities), *mass and chain retail*, and a *non-local* placebo of office-based activities such as consulting, legal services,

[13] The same degeneracy affects the commercial extensive margin when split. Because "thin" commercial blocks are those with no pre-shock transactions at all, their participation rate is bounded below at zero and the estimated coefficient is mechanically positive. I do not report those cells.

[14] The SII registry provides annual firm-level records including legal type, registered economic activity, and reported employment. I geocode firms to census blocks using their registered address observed before October 2019, so the treatment cohort is fixed on pre-shock information and cannot be contaminated by post-shock relocation or entry. Appendix 9 describes the geocoding procedure, the sample construction, and the activity classification.

Table 10: Market thickness and the margin of adjustment, post-normalization period

	(1) Thin blocks	(2) Thick blocks	(3) Thick − Thin
Panel A. Residential			
Sale price	−0.003	−0.057***	−0.054*
	(0.027)	(0.014)	(0.028)
Extensive margin	−0.0014	0.0048	0.0063
	(0.0024)	(0.0042)	(0.0047)
Panel B. Commercial (no interior thickness variation; pooled)			
Sale price		0.008 (0.046)	
Extensive margin		−0.0078*** (0.0027)	
Block FE	Yes	Yes	
Year-month (quarter) FE	Yes	Yes	
Municipality × year FE	Yes	Yes	

Notes. All entries are post-normalization effects of severe violence exposure, $\mathbb{1}\{V_b \geq 2\}$. The sale-price rows report effects on ln(price in UF) from transaction-level regressions; the extensive-margin rows report effects on an indicator for whether a block records any transaction in a given quarter. Market thickness is defined using pre-shock transaction counts through September 2019: "thick" blocks are those above the median among blocks with at least one sale, "thin" blocks those below it. Estimates in Panel A come from regressions run separately by cell; column (3) reports the difference and its standard error. Panel B reports pooled estimates because the median pre-shock commercial transaction rate across blocks is zero, leaving no interior variation on which to split; the forced thin price cell contains only 4,013 observations against 26,883 in the thick cell and is not identified. All specifications include property controls where applicable and the baseline fixed effects. Standard errors clustered at the block level.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

and investment vehicles that should be insensitive to street-level conditions.[15]

Employment in these data is attributed to the registered address rather than to the physical site of activity, and it is highly concentrated, so a single large employer in a block dominates the block total. The count of establishments reporting at least one worker is less exposed to both problems, since it weights every employer equally regardless of size. I therefore treat it as the primary outcome and report the worker-based specifications alongside it, so that the reader can see what each margin does and does not support.

Table 11 organizes the evidence by outcome and by firm group, and reports the pre-trend test for each. Three findings stand out.

First, firms do not die. The exit rate in directly local activities is a precise zero in the post-normalization period (−0.11 percentage points) and its pre-trend test is the cleanest in the table. Whatever violence did to local commerce, it did not push incumbent businesses out of the registry.

Second, firms shed their payrolls. The number of establishments reporting at least one worker falls by 13.5% in directly local activities, while the count of establishments merely present in the registry is unchanged (−1.9%). The gap between these two outcomes suggest that the same businesses remain registered but reduces employment. The decline is concentrated in specialized retail (−14.1%), the largest of the three subgroups with 25,215 firms, while restaurants and personal services move in the same direction without statistical significance. Upstream suppliers show no effect, and the non-local placebo is a precise zero, so the response is specific to activities that depend on foot traffic rather than a general shock to registered business in exposed blocks.

Third, the worker-count evidence points in the same direction but is weaker, and I report it as such. Total reported workers in directly local activities decline by 7.7% in the log specification and by 19.7% in Poisson, but neither is significant at conventional levels ($p = 0.11$ and $p = 0.38$). The employment margin that the data speak to clearly is therefore the number of establishments carrying a payroll, not the size of the payroll itself.

Figure 7 plots the event-study path for the establishment-count outcome. The post-shock decline is monotone and cumulative: near zero in 2019, −6% by 2020, and −14% by 2024, with

[15] Public entities, nonprofits, international organizations, and financial institutions are excluded from the universe. Online retailers are treated as ambiguous and excluded from all groups. The mass-and-chain group contains only 494 block–year cells and is reported for completeness but is too small to be informative. Appendix Table A.2 reports the full mapping.

confidence intervals excluding zero from 2020 onward. This is not a level shift at the moment of the unrest but a progressive attrition of employers over five years. The pre-period is less reassuring. Coefficients drift down to -3.7% by 2017 before returning to the 2018 reference, and the joint pre-trend test only passes marginally ($p = 0.087$). The corresponding test fails outright for the count of registered firms, which is a further reason to prefer the employer-count outcome, but it does not fully clear the establishment-count series either.

Read alongside Section 5.4, the two sources of commercial evidence agree on the same picture. Fewer commercial properties change hands, and the businesses occupying them operate at smaller scale without disappearing. Neither margin shows up in observed commercial prices, which is consistent with the thinness argument of Section 6.1: where transactions are rare, the shock is revealed in whether trade and employment occur at all rather than in the price at which trade clears. This is also the sense in which price-based estimates of the cost of violence are incomplete for commercial real estate.

6.3 Loss of safety as the underlying mechanism [still in process]

The evidence so far establishes that severely exposed blocks repriced, lost residents, and saw local commerce shed employment, and that none of this is attributable to the transit disruption. It does not yet say what buyers, renters, and business owners were responding to. A natural candidate is safety, if the unrest left a durable deterioration in local security conditions, rather than only a memory of eleven months of disorder, then the persistence of the discount five years later has a concrete referent. This subsection provides a first test of that idea. It is preliminary and I flag it as such.

I use reported crime at the level of the *Plan Cuadrante*, the police-patrol unit used by Carabineros de Chile, assigned to a couple of census blocks and compared between 2018—the last full pre-shock year—and 2023. The outcome is the change in the crime rate between the two years. All specifications include comuna fixed effects, so the comparison is between exposed and unexposed blocks within the same municipality, and standard errors are clustered at the comuna level across the 49 municipalities in the sample.

The city as a whole became less safe over this period. The mean block-level crime rate rose from 17.2 to 24.9, an increase of about 45%. The question is whether severely exposed blocks

Table 11: Firm-level adjustment in severely exposed blocks, post-normalization period

	(1) Direct local consumption	(2) Pre-trend p (col. 1)	(3) Upstream local suppliers	(4) Non-local placebo
Panel A. Outcomes				
Establishments present	-0.019	0.014	-0.053	0.022
	(0.019)		(0.039)	(0.018)
Establishments with employees	-0.135^{***}	0.087	-0.035	0.003
	(0.027)		(0.067)	(0.014)
Exit rate	-0.001	0.690	-0.001	-0.000
	(0.003)		(0.007)	(0.002)
Workers (PPML)	-0.197	0.074	-0.191	—
	(0.226)		(0.160)	
Workers ($\log(1+x)$)	-0.077	0.110	-0.043	—
	(0.049)		(0.101)	
Panel B. Establishments with employees, by subgroup within direct local consumption				
	Estimate	Pre-trend p	Firms	
Specialized retail	-0.141^{***}	0.072	25,215	
	(0.031)			
Restaurants and bars	-0.042	0.770	7,258	
	(0.062)			
Personal services	-0.100	0.424	2,963	
	(0.095)			
Block FE	Yes		Yes	Yes
Comuna × year FE	Yes		Yes	Yes

Notes. All entries report the post-normalization effect of severe violence exposure ($V_b \geq 2$) relative to unexposed blocks, from a balanced block–year panel built on the cohort of firms geocoded before October 2019 and covering 2015–2024, with 2018 as the reference year and 2019 treated as a transition year. Count outcomes are estimated by Poisson pseudo-maximum-likelihood, so coefficients are proportional effects; the exit rate is a linear probability. "Establishments present" counts firms appearing in the registry; "establishments with employees" counts those reporting at least one worker. Column (2) reports the p-value of a joint Wald test that all pre-shock coefficients equal zero, for the direct-local group. Standard errors clustered at the block level in parentheses.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

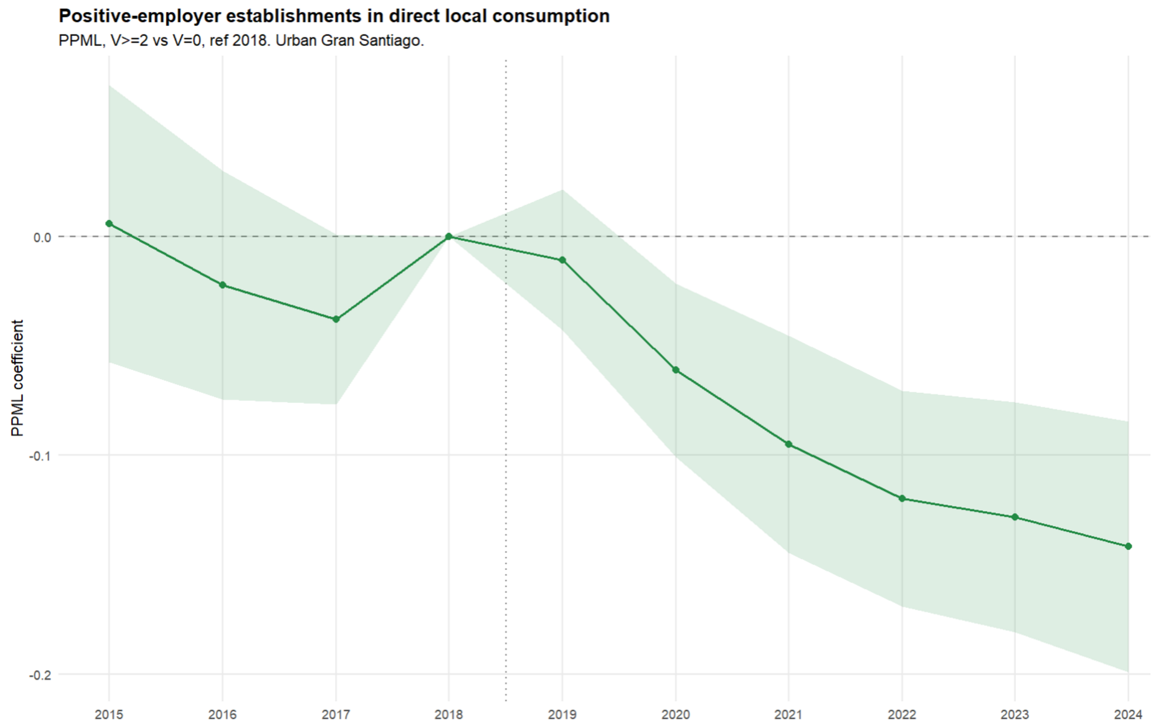

Figure 7: Event study: establishments with employees in directly local activities. Poisson pseudo-maximum-likelihood coefficients for severely exposed blocks ($V_b \geq 2$) relative to unexposed blocks ($V_b = 0$), with 2018 as the omitted reference year. The dependent variable is the count of establishments reporting at least one worker in directly local consumption activities, at the block–year level. The panel is balanced and built from the cohort of firms geocoded before October 2019. All models include block fixed effects and comuna-by-year fixed effects. The vertical dotted line marks the boundary between the pre-shock period and 2019, the transition year. Shaded bands indicate 95% confidence intervals based on standard errors clustered at the block level. The joint test that all pre-shock coefficients equal zero yields $p = 0.087$.

deteriorated faster than their neighbours, and Table 12 indicates that they did. Unconditionally, the crime rate rose by 6.97 points in unexposed blocks and by 17.30 in severely exposed ones. Within comuna, severe exposure is associated with an additional increase of 8.56 points. The effect is concentrated in robberies, which rise by an additional 3.01 points. Thefts show a large unconditional gap that disappears entirely once comuna fixed effects are included (0.38), indicating that the theft differential is a between-municipality pattern rather than a within-municipality one.

Three limitations keep this from being more than suggestive, and all three are structural rather than fixable by a different specification. There are only two time points, so there is no pre-trend to test and no way to distinguish a break at the unrest from a divergence already underway. Crime is measured at the resolution of the police cuadrante, which is coarser than the block, so exposed and unexposed blocks within the same cuadrante are assigned the same value and the estimate is attenuated toward zero by construction. And the 2018–2023 window contains the pandemic, a large migration episode, and a nationwide deterioration in security, any of which could interact with exposure. The distribution also matters, as the median change is 8.55 in exposed blocks against 6.18 in controls, a gap of 2.4 points rather than the 10.3 implied by the means, so the unconditional contrast is driven by the upper tail.

Table 12: Change in reported crime, 2018–2023, by violence exposure

	(1)	(2)	(3)
	Unconditional mean change		Within-comuna
	$V_b = 0$	$V_b \geq 2$	estimate
Total reported crime	6.97	17.30	8.561**
			(3.623)
Robberies	7.28	10.99	3.008*
			(1.766)
Thefts	1.15	4.11	0.381
			(3.004)
Comuna FE		Yes	
Blocks	40,076	3,765	43,841
Municipalities		49	

Notes. The dependent variable is the change in the reported crime rate between 2018 and 2023, measured at the level of the police *cuadrante* and assigned to census blocks. Columns (1) and (2) report raw mean changes by exposure group; column (3) reports the coefficient on severe exposure from a regression including comuna fixed effects, with standard errors clustered at the comuna level. Exposure is binary in this specification: severely exposed blocks ($V_b \geq 2$) against all others. Median changes are 6.18 and 8.55 for the two groups respectively, so the unconditional gap in means is influenced by the upper tail of the exposed distribution. Because crime is observed at cuadrante resolution, blocks within the same cuadrante share a value and the estimate is attenuated toward zero.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

7 Heterogeneous Incidence

Section 6 showed that the violence shock is revealed through different margins across markets: in residential real estate, primarily through prices; in commercial real estate, primarily through participation and operating scale. I now turn to a different question. Rather than asking *how* the shock appears in the data, this section asks *where* its effects are strongest. I focus on three pre-determined dimensions that are both empirically informative and conceptually relevant: prior crime exposure, urban location, and the socioeconomic level of the block.

In each case, I interact the severe violence indicator, $\mathbb{1}\{V_b \geq 2\}$, with a block characteristic measured before the shock and recover subgroup-specific post-normalization effects through linear combinations. Table 13 summarizes the results for residential sale prices, which is the outcome the paper identifies most cleanly. Full phase-by-phase estimates, and the corresponding splits for rents and commercial participation, are reported in the Appendix. One caution applies throughout: the three dimensions are correlated with one another, and each is estimated in a separate interaction rather than conditional on the others, so the panels should be read as three descriptions of the same underlying incidence and not as independent margins.

7.1 Prior crime exposure

I split blocks into high- and low-crime areas using the median 2018 robbery rate from the *Plan Cuadrante*. This is the sharpest contrast in the paper. For residential sale prices, the post-normalization effect of severe violence is small and statistically indistinguishable from zero in low-crime areas, at -1.8%, but reaches -6.1% in high-crime areas. The discount is therefore concentrated where the unrest reinforced an already adverse security environment, rather than where it introduced disorder into a previously safe one. This is also the heterogeneity result that connects most directly to the mechanism evidence of Section 6.3, since the same blocks that show the largest subsequent deterioration in reported crime are the ones where prior crime was already high.

The commercial participation margin does not reproduce this pattern with any precision. In low-crime blocks the post-normalization effect on the probability of any commercial transaction is -0.43 percentage points and insignificant. The commercial estimates are too imprecise across all subgroups to support a differential reading, which is consistent with the thinness argument of Section 6.1, since splitting an already sparse outcome by pre-determined characteristics leaves little identifying variation in each cell.

7.2 Urban location

I next classify blocks into core, pericentral, and peripheral areas according to their position in the metropolitan structure. The residential discount is strongly concentrated in the core, at -6.9%, against -4.3% in the periphery and an estimate close to zero and very imprecise in the pericentral ring. This is a change of emphasis relative to a reading in which the effect is spread evenly across the city: the shock hit hardest where residential value density is highest and where the disorder was most visible and most sustained.

Raw distance to the CBD does not reproduce this. Blocks within two kilometres of the centre show a post-normalization effect of −1.8%, indistinguishable from zero, and the distance bins display no monotone gradient. The contrast is informative rather than contradictory: what matters is the functional position of a block in the metropolitan structure—whether it belongs to the dense central fabric, the transitional ring, or the periphery—and not its Euclidean proximity to a point. It also provides a further argument against the centrality confound discussed in Section 5.3, since a mechanical repricing of central locations would show up in the distance bins.

Asking rents show their own spatial pattern, and it runs the other way. In the periphery the post-normalization rent effect is −7.6% and precisely estimated, far larger than the pooled estimate of −1.5% reported in Section 5.1. Peripheral exposed blocks therefore appear to adjust through rents in a way that central ones do not. I do not build on this, given that the rental series is used for corroboration rather than identification, but it suggests that the price–rent divergence documented earlier is itself spatially concentrated in the centre.

7.3 Socioeconomic level

The third dimension is the socioeconomic classification of the block, measured before the shock. The discount increases monotonically with socioeconomic level: −3.2% in low-status blocks, where it is imprecisely estimated, −4.5% in middle-status blocks, and −5.2% in high-status blocks. The gradient is the opposite of what a pure affordability story would predict, and the most natural reading is one of outside options. Buyers in higher-status segments can more readily relocate to an unaffected neighbourhood, so a deterioration in local security translates more fully into a price concession; buyers in lower-status segments face a narrower choice set and less scope to pay for safety, so the same deterioration is absorbed with a smaller discount.

This sits in apparent tension with the prior-crime result, since high-crime blocks are typically not high-status ones. The tension is more apparent than real: the two splits are correlated but far from collinear, particularly in the central comunas, where a dense apartment stock with mid-to-high socioeconomic classification coexists with high recorded crime. Because each dimension enters in a separate interaction, neither estimate is purged of the other, and I do not attempt to decompose them further.

Overall, the heterogeneity results sharpen the incidence of the shock without changing the message of the paper. The residential discount concentrates in central, higher-status blocks with pre-existing insecurity: places with more value at stake, residents with more alternatives, and a security problem that the unrest compounded rather than created. The commercial side remains too imprecise to support a differential reading, which is itself consistent with the asymmetry documented in Section 6: residential markets reveal violence through valuation, and commercial markets reveal it through whether trade occurs at all.

Table 13: Heterogeneous incidence of severe violence on residential sale prices, post-normalization period

	(1)	(2)	(3)
Panel A. Prior crime exposure (2018)			
	Low crime	High crime	
	−0.018	−0.061***	
	(0.019)	(0.016)	
Panel B. Urban form			
	Periphery	Pericentral	Core
	−0.043**	0.016	−0.069***
	(0.018)	(0.054)	(0.020)
Panel C. Socioeconomic level			
	Low	Middle	High
	−0.032	−0.045**	−0.052***
	(0.029)	(0.018)	(0.018)
Memo. Distance to the CBD			
	0–2 km		
	−0.018		
	(0.029)		
Block FE		Yes	
Year-month FE		Yes	
Municipality × year FE		Yes	

Notes. Each cell reports the post-normalization effect of severe violence, $\mathbb{1}\{V_b \geq 2\}$, on ln(price in UF) for the indicated subgroup, recovered as a linear combination of the main treatment effect and the relevant interaction term. Each panel comes from a separate regression interacting exposure with one pre-determined block characteristic; the panels are therefore not mutually conditioned. Prior crime exposure is defined using the median 2018 robbery rate from the *Plan Cuadrante*. Urban form classifies blocks by their position in the metropolitan structure. Socioeconomic level uses the block's pre-shock socioeconomic classification. All specifications include property controls and block, year-month, and municipality-by-year fixed effects. Corresponding estimates for asking rents and for commercial transaction participation, together with the full phase-by-phase results, are reported in Appendix Table **??**. Standard errors clustered at the block level.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

8 Robustness

This section evaluates the stability of the main residential sale-price result. Table A.16 summarizes the post-normalization coefficient on $\mathbb{1}\{V_b \geq 2\}$ across alternative specifications. The baseline estimate is a residential sale-price discount of approximately 5.6%. Full phase-by-phase results, together with additional falsification exercises for rents and commercial outcomes, are reported in the Appendix.

Alternative fixed effects and composition controls. The baseline specification includes block, year-month, and municipality-by-year fixed effects. Removing municipality-level time controls yields a much larger estimate of −10.0%, which suggests that part of the raw divergence reflects broader municipality-level shocks rather than violence exposure alone. Among the specifications that continue to absorb local time-varying heterogeneity, however, the estimate is stable. Replacing municipality-by-year with municipality-by-year-month fixed effects produces a nearly identical coefficient of −5.9%, adding property-type-by-year fixed effects gives −5.5%, and the accessibility-quintile controls discussed in Section 5.1 give −4.7%. Trimming rules make little difference: the estimate ranges from −5.6% to −5.3% as the sample is winsorized at successively more restrictive percentiles. A more demanding check replaces block fixed effects with property fixed effects, identifying the effect only from properties that transact more than once. On the 62,897 transactions belonging to 27,807 repeat-sale properties, the estimate is −5.9% ($p < 0.01$), so the discount is not driven by changes in the composition of the transacted stock within blocks.

Local boundary comparisons. To reduce concern that the baseline compares severely exposed blocks to distant and structurally different controls, I re-estimate the model within progressively narrower distances of the border between exposed and unexposed areas. The post-normalization coefficient is negative and significant at the 1% level at every bandwidth, ranging from −5.4% at 2,500m to −4.4% at 1,000m, where the control sample has contracted to half its original size. The mirror-image exercise, which drops rather than retains observations near the border, moves the estimate in the opposite direction, to −6.7% beyond 250m. Both are discussed in Section 5.2. Taken together they indicate that the discount is a localized phenomenon and that the baseline, if anything, understates it.

Pre-trends and temporal falsification. Pre-treatment coefficients for residential prices are jointly insignificant at every frequency the data support: $p = 0.958$ in the semiannual event study of Section 5.1, $p = 0.252$ quarterly, and $p = 0.068$ monthly. Because a joint Wald test cannot separate a systematic pre-trend from a few noisy periods, and only the former biases the estimator, I also test directly for a linear pre-shock divergence. The estimated pre-period slope for severely exposed blocks is -0.0007 log points per quarter ($p = 0.80$), implying a cumulative divergence of 0.6 log points over the two pre-shock years against a post-normalization effect of 5.6 log points. Low-exposure blocks, rents, and commercial prices are similarly flat, with no slope significant at any conventional level. There is no ramp for the design to extrapolate.[16]

Two falsification exercises complete the picture. Assigning a fake shock in April 2018 and re-estimating on pre-shock data only yields a precise zero for severe exposure in residential prices (-0.001, standard error 0.011), and small insignificant coefficients for rents (-0.005) and commercial prices ($+0.032$). A rolling placebo battery assigning five fake shock dates with symmetric twelve-month windows generates no residential coefficient as extreme as the actual estimate; the same exercise produces one placebo out of five for rents and five out of five for commercial prices, a further reason to treat the commercial price series as uninformative rather than as a null. The triple-difference specification of Section 5.3 also has jointly insignificant pre-treatment interaction terms ($p = 0.624$).

Spatial inference. The residential discount is robust to alternative inference procedures. Conley spatial standard errors with cutoffs of 1, 2, and 5 kilometres yield 0.015, 0.019, and 0.015 against a block-clustered 0.013, so the coefficient remains significant at the 1% level under every cutoff. Dropping one municipality at a time leaves the estimate in a range of -6.0% to -3.9%, always significant at the 1% level; the most influential single municipality is Santiago, whose exclusion moves the estimate to -3.9%. I also implement randomization inference by permuting treatment assignments within comuna-by-CBD strata, preserving the broad spatial structure of exposure while breaking the observed link between violence and outcomes. Across 500 permutations, only 0.2% of

[16] The joint tests are less stable across frequencies than the slopes, in both directions: the monthly test for rents passes at $p = 0.111$ while the quarterly test is marginal at $p = 0.050$, and the residential ordering runs the other way. This is what one expects when block-period cells are sparse. In the pre-shock period only 10.4% of residential block-months and 23.4% of block-quarters record any transaction, so a joint test over twenty-three monthly coefficients is dominated by sampling variation in thin cells. I therefore report all frequencies rather than designating one as canonical, and rely on the slope test as the diagnostic that maps onto the identifying assumption.

assignments produce a post-normalization coefficient at least as extreme as the observed estimate. This does not prove random assignment, but it makes it unlikely that the baseline result arises from broad spatial sorting alone.

Metro-threshold sensitivity. The channel results do not depend on the 500m cutoff used to define proximity to a closed station. Across thresholds from 300m to 800m (Appendix Table A.27), the *Violence Only* effect on residential prices stays between -5.0% and -5.5% and is significant at the 1% level in every case. The *Metro Only* coefficient is positive throughout and largest at the tightest cutoff, at $+7.2\%$ and significant at the 5% level for blocks within 300m of a closed station, decaying to zero by 800m. Blocks that lost station access and recorded no violence therefore trade at a premium, not a discount, and the premium is strongest precisely where the accessibility loss was most acute. The incremental effect of metro proximity among severely exposed blocks is never statistically distinguishable from zero (p between 0.23 and 0.97), while the severity gradient between $V_b \geq 2$ and $V_b = 1$ blocks far from closed stations is significant at the 1% level at every threshold. Commercial estimates remain imprecise throughout.

Predetermined exposure. The prosecutorial source used to construct violence exposure is based on spatial hotspots rather than fully dated incident-level records. This prevents a clean reconstruction of a predetermined exposure measure using only the first days or weeks of the unrest, so I cannot implement the strongest timing-based pre-determination check. This is a limitation of the data. At the same time, the unrest began abruptly and nationally on October 18, 2019, and the most intense violence was concentrated early in the episode, which limits the scope for exposure to be shaped by pre-existing property-market trends over a long horizon.

Taken together, these checks point in the same direction. The magnitude of the residential discount varies with the specification, as one would expect in a setting with overlapping shocks and spatial heterogeneity, and the monthly pre-trend test is a genuine caveat rather than a formality. But the sign, the timing, and the interpretation are stable across every variation considered, and the specifications that discipline the comparison most tightly—narrow boundary bandwidths, repeat sales, and the removal of contaminated controls—deliver estimates at or below the baseline rather than above it.

9 Conclusion

This paper studies how the October 2019 social unrest affected residential and commercial real estate markets in Santiago. Using georeferenced prosecutorial hotspot data linked to property transactions, rental listings, firm registries, and census records, I show that the same localized violence shock propagated through different margins across markets, and that it left a durable mark on valuations several years after the disorder itself had ended.

The first main finding is a persistent residential revaluation. Blocks with severe exposure to violence trade at a discount of roughly 5.6% in the post-normalization period, when metro service had resumed and pandemic restrictions were no longer binding. The effect is not a level shift that decays: it is smallest during the eleven months of transit disruption and largest three to six years after the shock. Blocks with low exposure show no comparable response at any horizon, a threshold pattern consistent with the idea that only more salient concentrations of violence generate a durable reassessment of neighborhood quality. Applied to the roughly 299,000 residential units located in severely exposed blocks, the discount implies a capitalized revaluation of about 42 million UF, or between USD 1.6 and 1.7 billion.

The second main finding is that this effect is attributable to violence rather than to the accompanying accessibility shock. Blocks near damaged stations but without recorded violence show no persistent discount, and if anything trade slightly above unexposed controls. Blocks with severe violence and no network disruption show the full effect. In the triple-difference specification, the interaction between violence and accessibility loss is positive rather than negative, so among exposed blocks it is those that lost *less* accessibility that display the larger discount. No weighting of a pure accessibility channel reproduces that pattern.

The third main finding concerns the relationship between prices and rents. Asking rents in severely exposed blocks fall during the disruption and the pandemic and then recover most of that decline, while sale prices continue to fall. Prices and rents therefore diverge, and the divergence is measurable: the implied gross capitalization rate on exposed residential property rises by about 4.8 log points, equivalent to roughly 0.3 percentage points on a baseline yield of 6.2%, and the estimate is stable across specification families and spatial units. The wedge is consistent with a higher required return on locations perceived as riskier and equally with lower expected growth in local rents, and the two are not separately identified here. The evidence on population loss and commercial retrenchment suggests that the second channel is at least partly active.

The fourth main finding is that the two markets adjust through different margins. In residential real estate, participation does not contract; if anything, severely exposed blocks transact more frequently after the shock, which points to rotation rather than to a frozen market and implies that the discount is realized rather than merely posted. In commercial real estate, I find no stable price response but a persistent decline of about 36% in the probability that a block records any commercial transaction. Firm-level records corroborate this without contradicting it: incumbent businesses do not exit the registry, but the number of establishments reporting any employee falls by around 13% in activities dependent on street-level demand. The contrast is economically informative. Where trading is thick enough for price discovery, the shock is capitalized into prices; where it is thin, the same shock is revealed through whether trade and employment occur at all. Only 24% of severely exposed blocks record a commercial transaction over the sample period, against 87% for residential.

The fifth main finding is that the revaluation has a real counterpart and an interpretable incidence. Severely exposed blocks lose about 3% of their population between the 2017 and 2024 censuses, so the price effect is not a paper loss without a corresponding reallocation of residents. The discount is concentrated in central blocks, in blocks of higher socioeconomic classification, and in blocks that already had high recorded crime before the unrest: places with more value at stake, residents with more alternatives, and a security problem that the episode compounded rather than created. Preliminary evidence on reported crime is consistent with that last reading, since severely exposed blocks show a larger subsequent deterioration in security than their neighbours within the same municipality.

More broadly, the paper contributes to the literature on conflict, crime, and urban real estate by examining a form of shock that is sudden, spatially concentrated, and neither a conventional war setting nor a single localized crime event. Episodic urban violence can leave persistent scars on local valuations, but those scars are not visible in the same way across asset markets, and estimating them from prices alone will understate them wherever the market is thin. The cap-rate result adds a further caution: even within a single market, a discount measured on sale prices and one measured on rents need not agree, and the gap between them is itself an object of interest rather than a measurement problem.

Several limitations remain. The violence measure is built from prosecutorial hotspots rather than fully dated incident-level records, which limits within-category variation and precludes a stronger predetermined-exposure design.Pre-treatment dynamics are flat in level and in slope, but the pre-shock window spans only three years, so the design cannot rule out divergences that began earlier. The rental series consists of posted asking rents rather than transacted contract rents, so landlords may adjust along margins the data do not record. Commercial valuation effects remain difficult to estimate precisely because the market becomes thin precisely where the shock is strongest. The firm-level and crime evidence is corroborative rather than identified: the former rests on a marginally passing pre-trend, and the latter on two cross-sections at police-district resolution, in a window that also contains the pandemic and a large migration episode. Future work could extend the analysis with incident-level timing, block-resolution crime at higher frequency, vacancy and contract-rent data, and a general-equilibrium treatment of re-sorting across locations.

References

Abadie, Alberto and Sofia Dermisi, "Is terrorism eroding agglomeration economies in Central Business Districts? Lessons from the office real estate market in downtown Chicago," *Journal of Urban Economics*, 9 2008, *64* (2), 451–463.

Araya, Martín Rosas, "Calculating Socio Material Territorial Index [R package ismtchile version 2.1.5]," *CRAN: Contributed Packages*, 7 2023.

Avendaño, Octavio, "ESTALLIDO SOCIAL EN CHILE: LOS DILEMAS POLÍTICOS DESDE OCTUBRE DEL 2019," *Política. Revista de Ciencia Política*, 12 2019, *57* (2), 105–119.

Besley, Timothy and Hannes Mueller, "Estimating the Peace Dividend: The Impact of Violence on House Prices in Northern Ireland," *American Economic Review*, 2012, *102* (2), 810–33.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, "Salience Theory of Choice Under Risk," 9 2010.

COES, "Protest Events Dataset 2008-2020," 2020.

Collins, William J and Robert A Margo, "The Economic Aftermath of the 1960s Riots in American Cities: Evidence from Property Values," *Journal of Economic History*, 12 2007, (67), 849–883.

Cox, Loreto, Ricardo González, and Carmen Le Foulon, "The 2019 Chilean Social Upheaval: A Descriptive Approach," *Journal of Politics in Latin America*, 4 2024, *16* (1), 68–89.

Diario Financiero, "Los impactos urbanos e inmobiliarios en el epicentro del estallido social — Diario Financiero."

EMOL, "El éxodo tras 18-O: Cerca de 900 comercios cerrados en el centro de Santiago y vacancia de oficinas disparada."

Fiscalia Nacional, "Informe Estallido Social," Technical Report, Ministerio Publico 2025.

Fisher, Jeffrey, Dean Gatzlaff, David Geltner, and Donald Haurin, "Controlling for the impact of variable liquidity in commercial real estate price indices," *Real Estate Economics*, 6 2003, *31* (2), 269–303.

Gibbons, Steve, "The costs of urban property crime," *Economic Journal*, 11 2004, *114* (499), F441–F463.

Gonzalez, Ricardo and Carmen Le Foulon Morán, "The 2019–2020 Chilean protests: A first look at their causes and participants," *International Journal of Sociology*, 5 2020, *50* (3), 227–235.

Gourley, Patrick, "Social Stigma and Asset Value," *Southern Economic Journal*, 1 2019, *85* (3), 919–938.

Haig, Christian Stirling, Katherine Schmidt, and Samuel Brannen, "The Age of Mass Protests: Understanding an Escalating Global Trend," 2020.

Hanushek, Eric A and John M Quigley, "The Dynamics of the Housing Market: A Stock Adjustment Model of Housing Consumption," *JOURNAL OF URBAN ECONOMICS*, 1979, *6*, 90–111.

Krainer, J., "A Theory of Liquidity in Residential Real Estate Markets," *Journal of Urban Economics*, 1 2001, *49* (1), 32–53.

La Tercera, "Vacancia de locales en el centro de Santiago disminuye en 2025: sector de Santa Lucía muestra el mejor desempeño - La Tercera."

Lens, Michael C and Rachel Meltzer, "IS CRIME BAD FOR BUSINESS? CRIME AND COMMERCIAL PROPERTY VALUES IN NEW YORK CITY *," *JOURNAL OF REGIONAL SCIENCE*, 2016, *56* (3), 442–470.

Linden, Leigh and Jonah E. Rockoff, "Estimates of the Impact of Crime Risk on Property Values from Megan's Laws," *American Economic Review*, 2008, *98* (3), 1103–27.

Metro de Santiago, "Memoria Anual," 2020.

Ooi, Joseph T.L. and Thao T.T. Le, "The spillover effects of infill developments on local housing prices," *Regional Science and Urban Economics*, 11 2013, *43* (6), 850–861.

Rosenthal, Stuart S. and Amanda Ross, "Violent crime, entrepreneurship, and cities," *Journal of Urban Economics*, 1 2010, *67* (1), 135–149.

Sampson, Robert J. and Stephen W. Raudenbush, "Systematic social observation of public spaces: A new look at disorder in urban neighborhoods," *American Journal of Sociology*, 1999, *105* (3), 603–651.

Sasse, Lea, "Chile Despertó-The Reasons for the Mass Protests in Chile Contact," 2019.

Wheaton, William C., "Vacancy, Search, and Prices in a Housing Market Matching Model," *https://doi.org/10.1086/261734*, 12 1990, *98* (6), 1270–1292.

Zahirovich-Herbert, Velma and Karen M. Gibler, "The effect of new residential construction on housing prices," *Journal of Housing Economics*, 2014, *26*, 1–18.

Appendix

Firm-Level Data: Construction and Classification

This appendix describes the construction of the firm-level panel used in Section 6.2. All data come from the publicly available corporate registry maintained by Chile's *Servicio de Impuestos Internos* (SII), which publishes annual snapshots of legally registered entities.[17]

Geocoding

The SII registry records each firm's registered headquarters address as a municipality code (*código de comuna*) and block code (*código de manzana*), which together identify a unique administrative unit within the SII's internal cadastral system. I match these to census-block identifiers (MANZENT_I) in two steps. First, I build a bridge between SII administrative codes (*cod_com*, *cod_mz*) and census-block identifiers using the property-level cadastral file (*propiedades Santiago urbano*), which contains both coding systems for all registered properties in the Santiago metropolitan region. I take the centroid of each SII administrative block and spatially join it to the 2017 census manzana shapefile. The bridge covers 97.1% of SII administrative blocks containing firms.

Second, I assign each firm to the census block corresponding to its registered headquarters address as of the last observation before October 2019. Firms that enter the registry after the shock are not geocoded and are therefore excluded from the analysis. This restriction is intentional: because addresses of post-shock entrants cannot be observed before treatment, including them would conflate entry decisions with treatment exposure. The analysis therefore covers the *pre-shock geocoded cohort* of firms, not the full universe of business activity. Of the approximately 1,009,000 unique firms observed at least once in the Santiago metropolitan region between 2015 and 2024, 658,000 (65.2%) are successfully matched to a census block. The remaining firms are registered outside the urban boundary, lack a valid block code, or cannot be matched through the bridge.

A key concern is whether match rates differ systematically by violence exposure. Table A.1

[17] The registry is published in two files: *Nómina de Personas Jurídicas* (PUB_NOMBRES_PJ), which records firm identity, legal type (*subtipo contribuyente*), and economic activity; and *Empresas por Tamaño según Ventas* (PUB_EMPRESAS_PJ), which records reported number of workers and sales bracket (*tramo de ventas*). I use annual snapshots from 2015 through 2024. A separate branch-level file provides establishment addresses for firms with multiple locations.

reports match rates by violence level. The share of firms successfully geocoded is virtually identical across exposure groups (65–67%), and a formal test of equality across violence categories fails to reject the null ($p = 0.41$). Differential selection into the geocoded sample is therefore unlikely to confound the difference-in-differences estimates.

Table A.1: Geocoding match rates by violence exposure

	$V_b = 0$	$V_b = 1$	$V_b \geq 2$
Firms in SII registry	632,000	198,000	179,000
Successfully geocoded	412,000	131,000	115,000
Match rate	65.2%	66.2%	64.2%

Notes. Counts are rounded to the nearest thousand. Match rates are computed as the share of unique firms (identified by RUT) observed in the Santiago metropolitan region between 2015 and 2024 that are successfully assigned to a census block via the SII-to-census bridge. Violence exposure is assigned using the block-level violence index V_b from the main analysis. Firms not matched to a census block are excluded from the firm-level analysis. A χ^2 test for equality of match rates across the three groups yields $p = 0.41$.

Panel construction

The geocoded firm sample yields 2,211,558 firm-year observations across 339,944 unique firms and 10 years (2015–2024). I observe each firm's reported number of workers and sales bracket (*tramo de ventas*) in each year it appears in the registry. The worker variable has complete coverage: 100% of firm-year observations report a value (including zero) for number of workers, with no differential coverage by violence exposure or calendar year.

For the mechanism analysis, I aggregate firm-level data to the block–year–group level, where group refers to the economic activity classification described below. The panel is balanced by construction using a cross-join of all block–group combinations that ever contain at least one firm of that group, crossed with all calendar years. Block–year–group cells with no active firms are assigned zero workers and zero firms, so that the outcome $\ln(1 + \text{total workers})$ equals zero rather than missing when a group empties out of a block. This balanced construction ensures that firm disappearance from a block is recorded as contraction rather than as a missing observation. The resulting panel contains 249,560 block–year–group observations, of which 34.1% have zero firms.

Activity classification

I classify firms in two stages. The first stage filters the institutional universe. I exclude entities whose legal type (*subtipo contribuyente*) indicates a public, nonprofit, or non-market entity: government agencies, municipalities, international organizations, unions, cooperatives, foundations, and other nonprofits (71,016 firm-years, 3.2% of the panel). I separately flag financial-sector entities—banks, pension fund administrators, insurance companies, and investment funds—as institutionally noisy, since their registered headquarters location is unlikely to reflect local economic exposure (3,436 firm-years).

The second stage classifies the remaining private commercial firms by registered economic activity (*actividad económica*) into three groups. Table A.2 reports the mapping. The *local-facing* group (253,140 firm-years) includes activities plausibly dependent on pedestrian traffic and neighborhood conditions: retail trade in food, clothing, and personal goods; restaurants, bars, and food service; personal services such as hairdressers and laundries; pharmacies and small medical centers; and vehicle repair. The *placebo* group (815,318 firm-years) includes activities plausibly less sensitive to street-level conditions: investment funds, management and IT consulting, legal and accounting services, engineering firms, real estate management, advertising, and administrative support services. Financial-sector entities flagged in the first stage are assigned to the placebo group. The remaining firms (1,072,084 firm-years) are classified as *ambiguous* and excluded from the main comparison.

The classification is intentionally broad and conservative. I do not attempt to identify individual firms that were directly affected by the unrest. Instead, the goal is to construct two groups with plausibly different average sensitivity to neighborhood-level disorder, so that a differential response can be interpreted as evidence of the commercial retrenchment channel.

Robustness of the firm-level results

Table A.3 reports the sensitivity of the main worker result to alternative samples and specifications. The baseline estimate from the balanced broad local-facing panel is a post-normalization decline of 7.4% ($p = 0.020$). Three alternative specifications preserve the negative sign but reduce precision.

First, a *strict* local-facing classification that excludes hardware stores, electronics retailers, vehicle repair, medical centers, gambling, and recreation activities reduces the sample from 253,140

Table A.2: Economic activity classification for firm-level mechanism analysis

Group	Activities included	Firm-years
Local-facing	Restaurants and mobile food service; bars and beverage service; catering; hotels, motels, and residentials; retail sale of food in small stores and minimarkets; retail sale of clothing, footwear, books, electronics, hardware, electrical appliances, pharmaceuticals, and cosmetics; specialized retail n.e.c.; hairdressing and beauty; laundry and dry cleaning; personal services n.e.c.; vehicle maintenance and repair; vehicle parts retail; dental and medical centers; gambling; sports and recreation activities	253,140
Placebo	Investment funds and financial holding entities; financial advisory; management consulting; IT consulting; engineering consulting; advertising services; legal services; accounting and auditing; purchase/sale/rental of real estate; real estate brokerage; rental of furnished properties; head offices; private security services; other business support services; banks, pension funds, insurance companies (institutional flag)	815,318
Ambiguous	Wholesale trade; construction; transportation and storage; manufacturing; education; large-format retail (supermarkets, department stores); fuel stations; online retail; health services not classified above; public administration residuals; "default value" entries; all other activities	1,072,084
Excluded	Government agencies; municipalities; international organizations; nonprofits; cooperatives; unions; foundations; community organizations; other entities without legal personality	71,016

Notes. Classification based on the registered economic activity (*actividad económica*) field in the SII corporate registry, supplemented by the legal type (*subtipo contribuyente*) for institutional filtering. Firm-year counts refer to the full 2015–2024 panel of geocoded firms. The local-facing and placebo groups are used in the main mechanism analysis (Table 11). Ambiguous activities are excluded from the local-versus-placebo comparison. A stricter local-facing definition that excludes hardware, electronics, vehicle-related, medical, gambling, and recreation activities (122,530 firm-years) is used as a robustness check.

to 122,530 local-facing firm-years and yields a post-normalization estimate of -4.4% ($t = -1.11$). The attenuation is consistent with lower statistical power rather than with misclassification, since the excluded activities are plausibly exposed to neighborhood conditions even if less clearly so than food service or personal care.

Second, restricting the sample to firms *without geocoded branch establishments* reduces the panel modestly (from 126,300 to 121,500 block-year observations) and yields -4.7% ($t = -1.55$). This check addresses the concern that multi-establishment firms may be assigned to headquarters locations that do not coincide with their actual areas of operation. The sign and magnitude are preserved, though statistical significance is lost.

Third, restricting to blocks with *at least three* local-facing firms in 2018 reduces the sample sharply (to 21,610 observations) and yields -4.4% ($t = -0.50$). The severe loss of power reflects the fact that most blocks have few local-facing firms, so that this restriction discards the majority of the identifying variation.

Across all specifications, the point estimates are negative and in the range of -4% to -7%, but only the baseline broad classification retains conventional statistical significance. I therefore interpret the firm-level evidence as suggestive of commercial retrenchment in exposed areas, consistent with the commercial transaction results, rather than as a precisely identified standalone finding.

Appendix: Additional Tables

Table A.3: Robustness of the firm-level worker result

	(1) Broad	(2) Strict	(3) No branches	(4) Stock ≥ 3
$V_b \geq 2$ × Post-normalization	-0.074^{**}	-0.044	-0.047	-0.044
	(0.032)	(0.039)	(0.030)	(0.087)
$V_b \geq 2$ × Metro closure	-0.063^{**}	-0.041	-0.041	-0.174^{**}
	(0.030)	(0.036)	(0.028)	(0.083)
$V_b \geq 2$ × COVID transition	-0.067^{**}	-0.068^{*}	-0.058^{*}	-0.148^{*}
	(0.032)	(0.038)	(0.030)	(0.089)
Pre-trend p-value	0.513	0.215	—	—
Block FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Comuna × year FE	Yes	Yes	Yes	Yes
Observations	126,300	84,580	121,500	21,610

Notes. All specifications use ln(1+total workers) at the block–year level as the dependent variable, estimated on the local-facing subsample of the balanced firm panel. Column (1) uses the broad local-facing classification. Column (2) restricts to a stricter set of activities excluding hardware, electronics, vehicle-related, medical, gambling, and recreation activities. Column (3) excludes firms with at least one geocoded branch establishment. Column (4) restricts to blocks with at least three local-facing firms in 2018. Standard errors clustered at the block level in parentheses. Pre-trend p-values are from joint Wald tests on pre-shock event-study coefficients for $V_b \geq 2$.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.4: Early protest geography and final violence exposure

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable	Viol. index	Viol. index	Viol. index	Any Viol.	Severe Viol.	Severe Viol.
Early KDE (z-score)	0.2717***	0.2004***	0.1334***	0.0995***	0.0582***	0.0370***
	(0.0031)	(0.0031)	(0.0350)	(0.0018)	(0.0012)	(0.0090)
Baseline controls	No	Yes	Yes	Yes	Yes	Yes
Municipality FE	No	No	Yes	No	No	Yes
Observations	48,803	48,803	48,803	48,803	48,803	48,803
R^2	0.1362	0.2375	0.4476	0.1944	0.1547	0.3136
Within R^2	–	–	0.2093	–	–	0.1425

This table provides a descriptive validation of the early geography of unrest using the strict early-exposure layer. The explanatory variable is a standardized kernel-density measure constructed from georeferenced protest points observed on October 18–19, 2019. The dependent variable is the final block-level violence classification built from Prosecutor hotspots observed between October 18, 2019 and March 31, 2020. Columns 1–3 use the discrete violence index $V_b \in \{0, 1, 2, 3\}$. Column 4 uses an indicator equal to one if $V_b > 0$. Columns 5–6 use an indicator equal to one if $V_b \geq 2$. Baseline controls include log distance to the CBD, log distance to the nearest closed metro station, log distance to the nearest metro station, log population density, log residential stock, and log commercial stock. Municipality FE models absorb municipality fixed effects. Standard errors in parentheses. This exercise is intended to show broad spatial alignment between early protest activity and the final hotspot geography, not to establish that final severe exposure was fully predetermined at the block level. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A.5: Baseline Difference-in-Differences Estimates, Residential Sale Prices

	(1) Cat.	(2) Cat. + Access	(3) Cont.	(4) Cont. + Access
$V_b = 1\times$ Metro closure	−0.0034	0.0028		
	(0.0120)	(0.0119)		
$V_b \geq 2\times$ Metro closure	−0.0269**	−0.0188		
	(0.0121)	(0.0121)		
$V_b = 1\times$ COVID transition	0.0013	0.0087		
	(0.0106)	(0.0107)		
$V_b \geq 2\times$ COVID transition	−0.0172	−0.0064		
	(0.0144)	(0.0142)		
$V_b = 1\times$ Post-normal	0.0018	0.0103		
	(0.0130)	(0.0132)		
$V_b \geq 2\times$ Post-normal	−0.0356***	−0.0231*		
	(0.0126)	(0.0129)		
$\tilde{V}_b^{std}\times$ Metro closure			−0.0084**	−0.0069**
			(0.0033)	(0.0032)
$\tilde{V}_b^{std}\times$ COVID transition			−0.0049	−0.0030
			(0.0032)	(0.0031)
$\tilde{V}_b^{std}\times$ Post-normal			−0.0105***	−0.0083**
			(0.0035)	(0.0035)
Block FE	Yes	Yes	Yes	Yes
Year-month FE	Yes	Yes	Yes	Yes
Municipality $\times$ year FE	Yes	Yes	Yes	Yes
Accessibility quintile $\times$ year FE	No	Yes	No	Yes
Observations	619,710	619,710	619,710	619,710
R^2	0.84426	0.84434	0.84426	0.84434
Within R^2	0.28389	0.28390	0.28388	0.28390

Notes: Dependent variable is ln(price in UF). Cat. uses the categorical exposure measure with $V_b = 0$ as the reference group. Cont. uses the standardized continuous violence intensity $\tilde{V}_b^{std}$. All specifications include property controls. Standard errors clustered at the block level.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.6: Baseline Difference-in-Differences Estimates, Asking Rents

	(1) Cat.	(2) Cat. + Access	(3) Cont.	(4) Cont. + Access
$V_b = 1\times$ Metro closure	-0.0286**	-0.0226**		
	(0.0114)	(0.0115)		
$V_b \geq 2\times$ Metro closure	-0.0274***	-0.0226**		
	(0.0106)	(0.0106)		
$V_b = 1\times$ COVID transition	-0.0272**	-0.0190*		
	(0.0114)	(0.0114)		
$V_b \geq 2\times$ COVID transition	-0.0319***	-0.0252**		
	(0.0099)	(0.0100)		
$V_b = 1\times$ Post-normal	-0.0094	-0.0012		
	(0.0120)	(0.0122)		
$V_b \geq 2\times$ Post-normal	-0.0369***	-0.0300***		
	(0.0110)	(0.0111)		
$\tilde{V}_b^{std}\times$ Metro closure			-0.0069***	-0.0062***
			(0.0023)	(0.0023)
$\tilde{V}_b^{std}\times$ COVID transition			-0.0079***	-0.0070***
			(0.0022)	(0.0022)
$\tilde{V}_b^{std}\times$ Post-normal			-0.0102***	-0.0092***
			(0.0024)	(0.0024)
Block FE	Yes	Yes	Yes	Yes
Year-month FE	Yes	Yes	Yes	Yes
Municipality $\times$ year FE	Yes	Yes	Yes	Yes
Accessibility quintile $\times$ year FE	No	Yes	No	Yes
Observations	43,044	42,715	43,044	42,715
R^2	0.61321	0.61235	0.61311	0.61228
Within R^2	0.26780	0.26816	0.26762	0.26802

Notes: Dependent variable is $\ln(\mathrm{UF/m^2})$. Cat. uses the categorical exposure measure with $V_b = 0$ as the reference group. Cont. uses the standardized continuous violence intensity $\tilde{V}_b^{std}$. All specifications include listing controls. Standard errors clustered at the block level.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.7: Baseline Difference-in-Differences Estimates, Commercial Sale Prices

	(1) Cat.	(2) Cat. + Access	(3) Cont.	(4) Cont. + Access
$V_b = 1\times$ Metro closure	-0.0117	-0.0234		
	(0.0716)	(0.0729)		
$V_b \geq 2\times$ Metro closure	0.0638	0.0585		
	(0.0510)	(0.0502)		
$V_b = 1\times$ COVID transition	0.0903*	0.0926		
	(0.0546)	(0.0582)		
$V_b \geq 2\times$ COVID transition	0.0661	0.0766		
	(0.0570)	(0.0610)		
$V_b = 1\times$ Post-normal	-0.0037	0.0169		
	(0.0563)	(0.0602)		
$V_b \geq 2\times$ Post-normal	0.0454	0.0691		
	(0.0590)	(0.0628)		
$\tilde{V}_b^{std}\times$ Metro closure			0.0079	0.0057
			(0.0108)	(0.0104)
$\tilde{V}_b^{std}\times$ COVID transition			0.0128	0.0139
			(0.0121)	(0.0127)
$\tilde{V}_b^{std}\times$ Post-normal			0.0048	0.0066
			(0.0123)	(0.0128)
Block FE	Yes	Yes	Yes	Yes
Year-month FE	Yes	Yes	Yes	Yes
Municipality $\times$ year FE	Yes	Yes	Yes	Yes
Accessibility quintile $\times$ year FE	No	Yes	No	Yes
Observations	31,023	29,295	31,023	29,295
R^2	0.84655	0.84914	0.84648	0.84907
Within R^2	0.62861	0.63683	0.62844	0.63666

Notes: Dependent variable is ln(price in UF). Cat. uses the categorical exposure measure with $V_b = 0$ as the reference group. Cont. uses the standardized continuous violence intensity $\tilde{V}_b^{std}$. All specifications include property controls. Standard errors clustered at the block level.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.8: Post-normalization effects by distance to nearest hotspot ($V_b = 0$ sample)

	Residential prices (1)	Asking rents (2)	Commercial prices (3)
0-250m × Post-normalization	-0.029	-0.044^{**}	0.042
	(0.023)	(0.022)	(0.092)
250-500m × Post-normalization	0.010	-0.016	0.108
	(0.020)	(0.022)	(0.108)
500-1,000m × Post-normalization	0.003	-0.018	0.009
	(0.012)	(0.018)	(0.082)
>1,000m		reference	
H_0: 0-250 = 500-1,000, p-value	0.183	0.245	0.748
Sample	$V_b = 0$	$V_b = 0$	$V_b = 0$
Block FE	Yes	Yes	Yes
Year-month FE	Yes	Yes	Yes
Municipality × year FE	Yes	Yes	Yes
Observations	396,053	21,549	13,147
R^2	0.869	0.630	0.887

Notes. Sample restricted to $V_b = 0$ blocks. Ring indicators are based on the distance from each block centroid to the nearest centroid of a block with $V_b \geq 1$, using the prosecutorial incident shapefile. The reference category is $V_b = 0$ blocks more than 1km from any hotspot. Only post-normalization coefficients are reported; earlier period interactions are estimated but suppressed. All specifications include the baseline controls and fixed effects. Standard errors are clustered at the block level. The bottom row reports the p-value for equality between the 0–250m and 500–1,000m coefficients.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.9: Channel Decomposition, Commercial Sale Prices

	Commercial Sale Prices
Metro closure period	
$V_b = 0$, near metro	0.076
	(0.156)
$V_b = 1$, far	0.004
	(0.074)
$V_b = 1$, near metro	−0.170
	(0.161)
$V_b \geq 2$, far	0.029
	(0.053)
$V_b \geq 2$, near metro	0.145**
	(0.063)
COVID transition	
$V_b = 0$, near metro	−0.062
	(0.143)
$V_b = 1$, far	0.099*
	(0.056)
$V_b = 1$, near metro	−0.103
	(0.140)
$V_b \geq 2$, far	0.065
	(0.064)
$V_b \geq 2$, near metro	0.048
	(0.069)
Post-normalization	
$V_b = 0$, near metro	0.001
	(0.102)
$V_b = 1$, far	0.009
	(0.057)
$V_b = 1$, near metro	−0.182
	(0.146)
$V_b \geq 2$, far	0.025
	(0.064)
$V_b \geq 2$, near metro	0.089
	(0.071)
Block FE	Yes
Year-month FE	Yes
Municipality × year FE	Yes
Observations	31,023
R^2	0.84665
Within R^2	0.62887

Notes: Six-group channel decomposition crossing violence exposure ($V_b = 0$, $V_b = 1$, $V_b \geq 2$) with metro proximity ($M_b = \mathbb{1}\{\text{dist. to closed station} \leq 500\text{m}\}$). Reference group is $V_b = 0$, far from closed stations, in the pre-shock period. Dependent variable is ln(price in UF). All specifications include commercial property controls, block fixed effects, year-month fixed effects, and municipality-by-year fixed effects. Standard errors clustered at the block level.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.10: Binary triple-difference estimates: all periods

	(1) Sale Prices	(2) Asking Rents
Metro closure period		
$V_b = 1$	-0.003	-0.027^{**}
	(0.012)	(0.012)
$V_b \geq 2$	-0.029^{**}	-0.032^{***}
	(0.013)	(0.011)
M_b	-0.004	0.014
	(0.019)	(0.015)
$V_b = 1 \times M_b$	-0.009	-0.042
	(0.043)	(0.027)
$V_b \geq 2 \times M_b$	0.016	0.022
	(0.028)	(0.021)
COVID transition		
$V_b = 1$	0.004	-0.027^{**}
	(0.011)	(0.012)
$V_b \geq 2$	-0.012	-0.029^{***}
	(0.015)	(0.010)
M_b	0.003	0.020
	(0.018)	(0.018)
$V_b = 1 \times M_b$	-0.032	-0.019
	(0.030)	(0.026)
$V_b \geq 2 \times M_b$	-0.038	-0.025
	(0.027)	(0.023)
Post-normalization		
$V_b = 1$	0.003	-0.007
	(0.013)	(0.012)
$V_b \geq 2$	-0.032^{**}	-0.038^{***}
	(0.013)	(0.012)
M_b	0.017	0.018
	(0.019)	(0.015)
$V_b = 1 \times M_b$	-0.014	-0.049^{*}
	(0.030)	(0.028)
$V_b \geq 2 \times M_b$	-0.036	-0.003
	(0.027)	(0.020)
Block FE	Yes	Yes
Year-month FE	Yes	Yes
Municipality × year FE	Yes	Yes
Observations	619,710	43,044
R^2	0.844	0.613

Notes. Triple-difference specification with $\mathbb{1}\{V_b = 1\}$, $\mathbb{1}\{V_b \geq 2\}$, binary metro proximity $M_b = \mathbb{1}\{\text{distance to a closed station} \leq 500\text{m}\}$, and their interactions by period. Reference group is $V_b = 0$, far from closed stations, in the pre-shock period. All specifications include property-level controls. Standard errors clustered at the block level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table A.11: Continuous triple-difference estimates: GTFS network accessibility shock

	(1) Sale Prices	(2) Asking Rents	(3) Commercial Prices
Metro closure period			
$\mathbb{1}\{V_b \geq 2\}$	-0.033^{***}	-0.018^{*}	0.040
	(0.011)	(0.009)	(0.046)
GTFS shock	-0.007	-0.014	-0.026
	(0.006)	(0.031)	(0.053)
$\mathbb{1}\{V_b \geq 2\}\times$ GTFS shock	0.024^{**}	0.095^{**}	0.157^{***}
	(0.009)	(0.042)	(0.041)
COVID transition			
$\mathbb{1}\{V_b \geq 2\}$	-0.022	-0.021^{**}	0.003
	(0.014)	(0.008)	(0.049)
GTFS shock	0.001	0.010	0.022
	(0.006)	(0.031)	(0.072)
$\mathbb{1}\{V_b \geq 2\}\times$ GTFS shock	0.011	0.056	0.162^{***}
	(0.010)	(0.037)	(0.042)
Post-normalization			
$\mathbb{1}\{V_b \geq 2\}$	-0.044^{***}	-0.034^{***}	0.028
	(0.012)	(0.009)	(0.049)
GTFS shock	0.013	0.003	0.121
	(0.008)	(0.031)	(0.092)
$\mathbb{1}\{V_b \geq 2\}\times$ GTFS shock	0.018^{*}	0.047	0.122^{**}
	(0.009)	(0.037)	(0.055)
Block FE	Yes	Yes	Yes
Year-month FE	Yes	Yes	Yes
Municipality $\times$ year FE	Yes	Yes	Yes
Observations	619,710	42,715	29,295
R^2	0.844	0.611	0.849

Notes. Continuous triple-difference specification replacing binary metro proximity with the GTFS-based network accessibility shock—the increase in travel time to the CBD when closed stations are removed from the transit network. The shock is scaled by its standard deviation without centering, so that zero corresponds to no network disruption. The main effect of $\mathbb{1}\{V_b \geq 2\}$ therefore captures the violence discount for blocks with zero accessibility disruption. All specifications include property-level controls. Standard errors clustered at the block level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table A.12: Continuous triple-difference estimates: station-distance disruption

	(1) Sale Prices	(2) Asking Rents	(3) Commercial Prices
Metro closure period			
$\mathbb{1}\{V_b \geq 2\}$	-0.033^{***}	-0.018^{*}	0.039
	(0.012)	(0.009)	(0.047)
Acc. shock (km)	-0.003	-0.013	-0.005
	(0.003)	(0.010)	(0.025)
$\mathbb{1}\{V_b \geq 2\}\times$ Acc. shock	0.007^{*}	0.051^{***}	0.056^{**}
	(0.004)	(0.018)	(0.026)
COVID transition			
$\mathbb{1}\{V_b \geq 2\}$	-0.021	-0.021^{**}	-0.003
	(0.014)	(0.008)	(0.050)
Acc. shock (km)	-0.003	-0.002	-0.001
	(0.005)	(0.011)	(0.032)
$\mathbb{1}\{V_b \geq 2\}\times$ Acc. shock	0.002	0.035^{**}	0.073^{***}
	(0.003)	(0.017)	(0.025)
Post-normalization			
$\mathbb{1}\{V_b \geq 2\}$	-0.041^{***}	-0.035^{***}	0.024
	(0.013)	(0.009)	(0.050)
Acc. shock (km)	0.004	-0.014	0.028
	(0.004)	(0.011)	(0.037)
$\mathbb{1}\{V_b \geq 2\}\times$ Acc. shock	0.004	0.036^{**}	0.047^{*}
	(0.003)	(0.017)	(0.024)
Block FE	Yes	Yes	Yes
Year-month FE	Yes	Yes	Yes
Municipality $\times$ year FE	Yes	Yes	Yes
Observations	619,710	42,715	29,295
R^2	0.844	0.611	0.849

Notes. Continuous triple-difference specification replacing binary metro proximity with the station-distance accessibility shock—the increase in Euclidean distance to the nearest metro station when closed stations are removed, measured in kilometers. The variable equals zero for blocks whose nearest station was not affected by closures. The main effect of $\mathbb{1}\{V_b \geq 2\}$ captures the violence discount for blocks with zero station-distance disruption. All specifications include property-level controls. Standard errors clustered at the block level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table A.13: Residential transaction activity

	(1) Categorical	(2) Categorical + R10	(3) Continuous	(4) Continuous + R10
Panel A. Extensive margin (N=1,704,456)				
$V_b = 1 \times$ Metro closure	−0.0001	−0.0004		
	(0.0025)	(0.0025)		
$V_b = 1 \times$ COVID transition	0.0018	0.0012		
	(0.0023)	(0.0023)		
$V_b = 1 \times$ Post-normalization	0.0040*	0.0029		
	(0.0022)	(0.0023)		
$V_b \geq 2 \times$ Metro closure	−0.0036	−0.0042		
	(0.0031)	(0.0032)		
$V_b \geq 2 \times$ COVID transition	−0.0009	−0.0024		
	(0.0027)	(0.0028)		
$V_b \geq 2 \times$ Post-normalization	0.0055**	0.0022		
	(0.0027)	(0.0028)		
$\tilde{V}_b^{std} \times$ Metro closure			−0.0042***	−0.0045***
			(0.0013)	(0.0014)
$\tilde{V}_b^{std} \times$ COVID transition			0.0005	−0.0001
			(0.0012)	(0.0013)
$\tilde{V}_b^{std} \times$ Post-normalization			0.0021	0.0007
			(0.0013)	(0.0014)
Panel B. Transaction counts, PPML (N=1,333,465)				
$V_b = 1 \times$ Metro closure	−0.0310	−0.0122		
	(0.1050)	(0.1033)		
$V_b = 1 \times$ COVID transition	0.0119	0.0297		
	(0.0850)	(0.0887)		
$V_b = 1 \times$ Post-normalization	0.1124	0.1134		
	(0.0897)	(0.0928)		
$V_b \geq 2 \times$ Metro closure	0.0791	0.0918		
	(0.1047)	(0.1075)		
$V_b \geq 2 \times$ COVID transition	0.1418	0.1586		
	(0.0920)	(0.0986)		
$V_b \geq 2 \times$ Post-normalization	0.1560	0.1492		
	(0.1076)	(0.1137)		
$\tilde{V}_b^{std} \times$ Metro closure			0.0372	0.0386
			(0.0335)	(0.0337)
$\tilde{V}_b^{std} \times$ COVID transition			0.0507*	0.0522*
			(0.0307)	(0.0315)
$\tilde{V}_b^{std} \times$ Post-normalization			0.0298	0.0261
			(0.0311)	(0.0318)
Panel C. Turnover rate(N=1,704,456)				
$V_b = 1 \times$ Metro closure	−0.0003	−0.0003		
	(0.0003)	(0.0003)		
$V_b = 1 \times$ COVID transition	0.0000	0.0002		
	(0.0003)	(0.0003)		
$V_b = 1 \times$ Post-normalization	0.0034	0.0037		
	(0.0027)	(0.0027)		
$V_b \geq 2 \times$ Metro closure	0.0000	0.0000		
	(0.0003)	(0.0003)		
$V_b \geq 2 \times$ COVID transition	0.0003	0.0004		
	(0.0004)	(0.0006)		
$V_b \geq 2 \times$ Post-normalization	0.0022*	0.0016		
	(0.0013)	(0.0016)		
$\tilde{V}_b^{std} \times$ Metro closure			−0.0002*	−0.0002*
			(0.0001)	(0.0001)
$\tilde{V}_b^{std} \times$ COVID transition			0.0000	0.0000
			(0.0002)	(0.0002)
$\tilde{V}_b^{std} \times$ Post-normalization			0.0011*	0.0008
			(0.0006)	(0.0006)
Block FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Municipality × year FE	Yes	Yes	Yes	Yes
Accessibility quintile × year FE	No	Yes	No	Yes

Notes. Panel A uses an indicator for any transaction in a block-quarter. Panel B uses PPML for transaction counts on the stock-positive universe; block fixed effects with only zero outcomes are dropped automatically by the estimator. Panel C uses the transaction-to-stock ratio. All specifications include block fixed effects, quarter fixed effects, and municipality-by-year fixed effects. Standard errors clustered at the block level.

*** p < 0.01, ** p < 0.05, * p < 0.10

Table A.14: Commercial transaction activity

	(1) Categorical	(2) Categorical + R10	(3) Continuous	(4) Continuous + R10
Panel A. Extensive margin (N=486,540)				
$V_b = 1 \times$ Metro closure	−0.0062**	−0.0062**		
	(0.0025)	(0.0025)		
$V_b = 1 \times$ COVID transition	−0.0018	−0.0015		
	(0.0025)	(0.0025)		
$V_b = 1 \times$ Post-normalization	−0.0020	−0.0019		
	(0.0022)	(0.0022)		
$V_b \geq 2 \times$ Metro closure	−0.0164***	−0.0165***		
	(0.0034)	(0.0034)		
$V_b \geq 2 \times$ COVID transition	−0.0073**	−0.0067**		
	(0.0030)	(0.0031)		
$V_b \geq 2 \times$ Post-normalization	−0.0078***	−0.0074***		
	(0.0027)	(0.0028)		
$\tilde{V}_b^{std} \times$ Metro closure			−0.0068***	−0.0068***
			(0.0014)	(0.0014)
$\tilde{V}_b^{std} \times$ COVID transition			−0.0032***	−0.0031**
			(0.0013)	(0.0013)
$\tilde{V}_b^{std} \times$ Post-normalization			−0.0034***	−0.0032***
			(0.0011)	(0.0012)
Panel B. Transaction counts, PPML (N=164,411)				
$V_b = 1 \times$ Metro closure	−0.3988**	−0.4398**		
	(0.1878)	(0.1924)		
$V_b = 1 \times$ COVID transition	−0.1250	−0.1031		
	(0.2035)	(0.2059)		
$V_b = 1 \times$ Post-normalization	−0.1387	−0.1705		
	(0.2021)	(0.2026)		
$V_b \geq 2 \times$ Metro closure	0.1510	0.1202		
	(0.1968)	(0.1999)		
$V_b \geq 2 \times$ COVID transition	0.0855	0.1014		
	(0.1455)	(0.1485)		
$V_b \geq 2 \times$ Post-normalization	0.2824*	0.2502		
	(0.1566)	(0.1590)		
$\tilde{V}_b^{std} \times$ Metro closure			0.0387	0.0350
			(0.0397)	(0.0403)
$\tilde{V}_b^{std} \times$ COVID transition			0.0279	0.0307
			(0.0409)	(0.0413)
$\tilde{V}_b^{std} \times$ Post-normalization			0.0804*	0.0770
			(0.0483)	(0.0487)
Panel C. Turnover rate (N=486,540)				
$V_b = 1 \times$ Metro closure	0.0003	0.0004		
	(0.0010)	(0.0010)		
$V_b = 1 \times$ COVID transition	0.0014	0.0014		
	(0.0010)	(0.0011)		
$V_b = 1 \times$ Post-normalization	0.0015	0.0016*		
	(0.0009)	(0.0010)		
$V_b \geq 2 \times$ Metro closure	0.0003	0.0003		
	(0.0010)	(0.0010)		
$V_b \geq 2 \times$ COVID transition	0.0010	0.0008		
	(0.0013)	(0.0015)		
$V_b \geq 2 \times$ Post-normalization	0.0016	0.0016		
	(0.0017)	(0.0017)		
$\tilde{V}_b^{std} \times$ Metro closure			0.0003	0.0003
			(0.0003)	(0.0003)
$\tilde{V}_b^{std} \times$ COVID transition			0.0003	0.0003
			(0.0003)	(0.0003)
$\tilde{V}_b^{std} \times$ Post-normalization			0.0001	0.0000
			(0.0003)	(0.0003)
Block FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Municipality × year FE	Yes	Yes	Yes	Yes
Accessibility quintile × year FE	No	Yes	No	Yes

Notes. Panel A uses an indicator for any transaction in a block-quarter. Panel B uses PPML for transaction counts on the stock-positive universe; block fixed effects with only zero outcomes are dropped automatically by the estimator. Panel C uses the transaction-to-stock ratio. All specifications include block fixed effects, quarter fixed effects, and municipality-by-year fixed effects. Standard errors clustered at the block level.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$

Table A.15: Triple-Difference Volume Regressions with Standardized Accessibility Shock

	(1) Res. Extensive	(2) Res. Counts (PPML)	(3) Com. Extensive	(4) Com. Counts (PPML)
treat_sev × Metro closure	−0.0057*	0.0717	−0.0162***	0.4177**
	(0.0034)	(0.0916)	(0.0037)	(0.2062)
treat_sev × COVID transition	−0.0007	0.0961	−0.0076**	0.1177
	(0.0029)	(0.0778)	(0.0032)	(0.1488)
treat_sev × Post-normalization	0.0047	0.0991	−0.0084***	0.3389**
	(0.0030)	(0.0926)	(0.0029)	(0.1690)
treat_sev × acc_shock_std × Metro closure	0.0045*	−0.0461	0.0041**	0.4361
	(0.0023)	(0.0724)	(0.0018)	(0.2707)
treat_sev × acc_shock_std × COVID transition	−0.0016	−0.1261**	0.0032**	−0.1086
	(0.0020)	(0.0562)	(0.0016)	(0.0997)
treat_sev × acc_shock_std × Post-normalization	−0.0005	−0.0329	0.0043***	−0.0608
	(0.0020)	(0.0636)	(0.0015)	(0.1081)
Block FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Municipality × year FE	Yes	Yes	Yes	Yes
Estimator	OLS	PPML	OLS	PPML
Observations	1,704,456	1,333,465	486,540	164,411
R^2 / Pseudo-R^2	0.30667	0.51544	0.28234	0.44079

Notes. This table reports triple-difference specifications for transaction volumes using the standardized accessibility shock measure, `acc_shock_std`. Columns (1) and (3) use the probability of any transaction as the dependent variable and are estimated by OLS. Columns (2) and (4) use transaction counts and are estimated by Poisson pseudo-maximum likelihood. All specifications include block fixed effects, quarter fixed effects, and municipality-by-year fixed effects. Standard errors are clustered at the block level.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.16: Robustness of the Post-Normalization Violence Effect on Residential Sale Prices

Specification	Coef.	SE	Sig.
Fixed effects and controls			
Baseline (block + YM + municipality $\times$ year)	-0.056	(0.013)	***
TWFE (block + YM only)	-0.100	(0.010)	***
Municipality $\times$ year-month	-0.059	(0.012)	***
+ Property type $\times$ year	-0.055	(0.013)	***
+ Accessibility quintile $\times$ year	-0.047	(0.013)	***
+ CBD bin $\times$ year	-0.042	(0.013)	***
Sample restrictions			
Boundary bandwidth 2,500m	-0.054	(0.013)	***
Boundary bandwidth 1,500m	-0.052	(0.013)	***
Boundary bandwidth 1,000m	-0.044	(0.015)	***
Donut, drop within 250m of edge	-0.067	(0.014)	***
Trim 1/99	-0.056	(0.013)	***
Trim 2/98	-0.054	(0.013)	***
Trim 5/95	-0.053	(0.012)	***
Repeat sales (property FE)	-0.059	(0.017)	***
Alternative inference			
Conley SE (1km)	-0.056	(0.015)	***
Conley SE (2km)	-0.056	(0.019)	***
Conley SE (5km)	-0.056	(0.015)	***
Memo			
Leave-one-municipality-out range	$[-0.060,\ -0.039]$, all $p < 0.01$		
Randomization inference (500 permutations)	$p = 0.002$		

Notes: Each row reports the post-normalization coefficient on $\mathbb{1}\{V_b \geq 2\}$ from a variant of the baseline residential sale-price specification. The baseline includes block, year-month, and municipality-by-year fixed effects with standard errors clustered at the block level. Boundary bandwidths restrict the sample to observations within the indicated distance of the exposed/unexposed border; the donut row drops observations within the indicated distance instead. Repeat sales replaces block fixed effects with property fixed effects, restricting the sample to the 62,897 transactions of properties that sell more than once. Conley standard errors use spherical distance. Randomization inference permutes treatment within comuna-by-CBD strata. Full phase-by-phase results and additional checks are reported in the Appendix.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.17: Alternative Fixed-Effects Structures (Residential Sale Prices)

	(1) TWFE	(2) Baseline	(3) Comuna × YM
$V_b = 1 \times$ Post-norm.	−0.029***	−0.005	−0.004
	(0.010)	(0.013)	(0.013)
$V_b \geq 2 \times$ Post-norm.	−0.100***	−0.056***	−0.059***
	(0.010)	(0.013)	(0.012)
Block FE	Yes	Yes	Yes
Year-month FE	Yes	Yes	—
Municipality × year	—	Yes	—
Municipality × year-month	—	—	Yes
Observations	619,710	619,710	619,710
R^2	0.843	0.844	0.847

Notes: Dependent variable: ln(price in UF). Column (2) is the baseline specification from Table 3. All columns include property controls. Standard errors clustered at the block level.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.18: Alternative Fixed-Effects Structures (Commercial Prices and Rents)

	Commercial Prices			Asking Rents		
	TWFE	Baseline	Com.×YM	TWFE	Baseline	Com.×YM
$V_b = 1 \times$ Post-norm.	−0.028	−0.042	−0.073	−0.002	0.014	0.015
	(0.048)	(0.055)	(0.060)	(0.012)	(0.010)	(0.011)
$V_b \geq 2 \times$ Post-norm.	−0.084**	−0.012	−0.049	−0.098***	−0.016*	−0.012
	(0.039)	(0.056)	(0.056)	(0.011)	(0.009)	(0.010)
Block FE	Yes	Yes	Yes	Yes	Yes	Yes
Year-month FE	Yes	Yes	—	Yes	Yes	—
Municipality × year	—	Yes	—	—	Yes	—
Municipality × year-month	—	—	Yes	—	—	Yes
Observations	31,023	31,023	31,023	674,479	674,479	674,479
R^2	0.842	0.846	0.866	0.613	0.635	0.638

Notes: Commercial dependent variable: ln(price in UF). Rent dependent variable: ln(UF/m^2). All columns include property controls. Standard errors clustered at the block level.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.19: Boundary Bandwidth Design

	1,000m	1,500m	2,000m	2,500m
Panel A. Residential sale prices				
$V_b = 1$ × Post-norm.	−0.003	−0.003	−0.003	−0.003
	(0.015)	(0.014)	(0.013)	(0.013)
$V_b \geq 2$ × Post-norm.	−0.044***	−0.052***	−0.051***	−0.054***
	(0.015)	(0.013)	(0.013)	(0.013)
Observations	305,242	408,499	479,211	517,399
R^2	0.791	0.801	0.809	0.815
Panel B. Commercial sale prices				
$V_b = 1$ × Post-norm.	−0.089	−0.103*	−0.095*	−0.043
	(0.065)	(0.062)	(0.058)	(0.057)
$V_b \geq 2$ × Post-norm.	−0.032	−0.058	−0.063	−0.020
	(0.065)	(0.062)	(0.059)	(0.059)
Observations	15,979	21,669	23,922	25,531
R^2	0.855	0.850	0.850	0.848

Notes: Sample restricted to observations within the indicated distance of the $V_b > 0$ / $V_b = 0$ boundary, measured using signed distance to the nearest hotspot polygon edge. The rental sample does not carry the signed-distance variable and is therefore omitted. All specifications include block, year-month, and municipality-by-year fixed effects plus property controls. Standard errors clustered at the block level.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.20: Pre-Trend Tests: Joint Wald Tests and Pre-Period Slopes

	Monthly		Quarterly		Semiannual	Pre-period slope	
Market	$V_b\!=\!1$	$V_b\!\geq\!2$	$V_b\!=\!1$	$V_b\!\geq\!2$	Violence	Estimate	p
Residential prices	0.443	0.068	0.343	0.252	0.958	−0.0007	0.804
Asking rents	0.771	0.111	0.673	0.050	0.260	0.0001	0.974
Commercial prices	0.132	0.101	0.947	0.523	0.649	−0.0032	0.879
Memo. Pre-shock cell density (share of block-periods with any transaction)							
Residential	10.4%		23.4%				
Asking rents	21.8%		38.9%				
Commercial	6.8%		16.4%				

Notes: The monthly and quarterly columns report p-values from joint Wald tests of the null that all pre-shock event-study interactions equal zero, using months −24 to −1 and quarters −8 to −2 respectively, each relative to the period immediately preceding the shock. The semiannual column tests the pre-shock violence interactions in the event-time specification of Section 5.1. The slope columns report the coefficient on the interaction between severe exposure and a linear quarterly time trend, estimated on pre-shock data only, in log points per quarter; the corresponding cumulative divergence over the two pre-shock years is −0.006 for residential prices, 0.001 for rents, and −0.026 for commercial prices. Quarterly specifications use calendar-quarter rather than year-month fixed effects. All specifications include property controls and the baseline fixed-effects structure for the corresponding market, with standard errors clustered at the block level.

Table A.21: Rolling Temporal Placebo (Pre-Shock Fake Dates)

	Actual estimate (symmetric window)	Placebo dates tested	Share of placebos at least as extreme
Residential sale prices	−0.035	5	0.00
Asking rents	−0.023	5	0.20
Commercial sale prices	0.010	5	1.00

Notes: Each row assigns five fake shock dates in the pre-shock period and re-estimates the baseline specification around each with a symmetric twelve-month window. The actual estimate is re-computed on the same symmetric window so that it is comparable to the placebo distribution, which is why it differs from the headline post-normalization coefficient. The final column reports the share of placebo coefficients at least as large in absolute value as the actual one. A share near zero indicates that the estimate is not reproducible at arbitrary pre-shock dates.

Table A.22: Composition Controls: Property Type × Year Fixed Effects

	Residential Prices		Commercial Prices
	Baseline	+ Type×Year	+ Type×Year
$V_b = 1$ × Post-norm.	−0.005	−0.004	−0.012
	(0.013)	(0.013)	(0.055)
$V_b \geq 2$ × Post-norm.	−0.056***	−0.055***	0.015
	(0.013)	(0.013)	(0.055)
Block FE	Yes	Yes	Yes
Year-month FE	Yes	Yes	Yes
Mun.×year FE	Yes	Yes	Yes
Type×year FE	—	Yes	Yes
Observations	619,710	619,710	31,023
R^2	0.844	0.844	0.847

Notes: Residential type is apartment vs. house (is_apt×year). Commercial type is office vs. retail (is_office × year). All specifications include property controls. Standard errors clustered at the block level.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.23: Winsorization Sensitivity

	Trim 1/99	Trim 2/98	Trim 5/95
Panel A. Residential sale prices			
$V_b = 1 \times$ Post-norm.	−0.002	0.001	0.009
	(0.013)	(0.014)	(0.014)
$V_b \geq 2 \times$ Post-norm.	−0.056***	−0.054***	−0.053***
	(0.013)	(0.013)	(0.012)
Observations	607,332	594,936	557,859
Panel B. Asking rents			
$V_b = 1 \times$ Post-norm.	0.009	0.007	0.002
	(0.009)	(0.009)	(0.008)
$V_b \geq 2 \times$ Post-norm.	−0.020**	−0.020**	−0.023***
	(0.008)	(0.008)	(0.007)
Observations	660,993	647,509	607,031
Panel C. Commercial sale prices			
$V_b = 1 \times$ Post-norm.	−0.033	−0.010	0.008
	(0.050)	(0.049)	(0.052)
$V_b \geq 2 \times$ Post-norm.	−0.002	−0.010	−0.017
	(0.054)	(0.055)	(0.056)
Observations	30,401	29,781	27,919

Notes: Sample trimmed at the indicated percentiles of the dependent variable. The baseline specification (Table 3) trims at 1/99 on the original price in UF before taking logs. All specifications include block, year-month, and municipality-by-year fixed effects plus property controls. Standard errors clustered at the block level.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.24: Repeat-Sales Specification: Property Fixed Effects

	Residential		Commercial
	Block FE	Property FE	Property FE
$V_b = 1 \times$ Post-norm.	−0.005	−0.028*	0.076
	(0.013)	(0.016)	(0.100)
$V_b \geq 2 \times$ Post-norm.	−0.056***	−0.059***	−0.062
	(0.013)	(0.017)	(0.090)
Block FE	Yes	—	—
Property FE	—	Yes	Yes
Year-month FE	Yes	Yes	Yes
Municipality × year FE	Yes	Yes	Yes
Properties	—	27,807	1,812
Observations	619,710	62,897	4,033
R^2	0.844	0.935	0.946

Notes: The repeat-sales columns replace block fixed effects with property fixed effects, so that identification comes only from properties transacting more than once during the sample period. Standard errors are two-way clustered by property and block. All specifications include property controls where not absorbed.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.25: Spatial Inference: Alternative Standard Errors

	Block	Conley 1km	Conley 2km	Conley 5km
Panel A. Residential sale prices				
$V_b = 1 \times$ Post-norm.	-0.005	-0.005	-0.005	-0.005
	(0.013)	(0.017)	(0.018)	(0.012)
$V_b \geq 2 \times$ Post-norm.	-0.056***	-0.056***	-0.056***	-0.056***
	(0.013)	(0.015)	(0.019)	(0.015)
Panel B. Commercial sale prices				
$V_b = 1 \times$ Post-norm.	-0.042	-0.042	-0.042	-0.042
	(0.055)	(0.061)	(0.027)	(0.053)
$V_b \geq 2 \times$ Post-norm.	-0.012	-0.012	-0.012	-0.012
	(0.056)	(0.049)	(0.048)	(0.023)
Observations	619,710 (Panel A); 31,023 (Panel B)			

Notes: All columns estimate the same baseline specification; only the variance estimator differs. Conley standard errors use spherical distance with the indicated cutoff radius.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.26: Transaction Quantity: Alternative Estimators

	(1) Extensive margin (any transaction)	(2) Counts (PPML)	(3) Counts, active blocks only
Panel A. Residential			
$V_b = 1 \times$ Post-norm.	0.004	0.134	0.132
	(0.002)	(0.090)	(0.103)
$V_b \geq 2 \times$ Post-norm.	0.006**	0.197*	0.195*
	(0.003)	(0.109)	(0.115)
Panel B. Commercial			
$V_b = 1 \times$ Post-norm.	-0.001*	-0.110	-0.182
	(0.001)	(0.201)	(0.239)
$V_b \geq 2 \times$ Post-norm.	-0.004***	0.279	0.279
	(0.001)	(0.156)	(0.183)

Notes: Column (1) is a linear probability model for whether a block records any transaction in a given quarter. Columns (2) and (3) estimate transaction counts by Poisson pseudo-maximum-likelihood; column (3) restricts the sample to block-quarters with positive activity (216,128 of 1,756,908 for residential and 13,786 for commercial). The contrast between the negative extensive-margin effect and the positive count effect in the commercial market is discussed in Section 5.4: participation falls while activity concentrates among the blocks that remain active. All specifications include block, quarter, and municipality-by-year fixed effects. Standard errors clustered at the block level.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.27: Metro Proximity Threshold Sensitivity: Residential Sale Prices

Threshold (m)	300	400	500	600	700	800
Post-normalization coefficients						
Metro Only ($V_b = 0$, near)	0.072**	0.031	0.031*	0.026	0.006	0.005
	(0.028)	(0.024)	(0.019)	(0.019)	(0.018)	(0.018)
$V_b = 1$, far	−0.003	−0.003	−0.002	−0.002	0.001	0.001
	(0.013)	(0.013)	(0.014)	(0.014)	(0.014)	(0.015)
$V_b \geq 2$, far	−0.053***	−0.055***	−0.050***	−0.050***	−0.051***	−0.053***
	(0.013)	(0.013)	(0.013)	(0.014)	(0.014)	(0.014)
$V_b \geq 2$, near metro	−0.069**	−0.054**	−0.073***	−0.065***	−0.064***	−0.061***
	(0.034)	(0.024)	(0.020)	(0.018)	(0.018)	(0.017)
p-values for formal contrasts						
Metro Only = 0	0.011	0.193	0.096	0.169	0.728	0.794
Violence Only = 0	0.000	0.000	0.000	0.000	0.000	0.000
Compounding	0.621	0.965	0.231	0.324	0.457	0.602
Severity gradient	0.004	0.003	0.007	0.007	0.003	0.003
Observations			619,710			

Notes: Six-group channel decomposition re-estimated at each metro proximity threshold. Metro Only = $V_b = 0$ within the threshold of a closed station; Violence Only = $V_b \geq 2$ beyond it. Compounding tests H_0: ($V_b \geq 2$, near) − ($V_b \geq 2$, far) = 0. Severity gradient tests H_0: ($V_b \geq 2$, far) − ($V_b = 1$, far) = 0. Contrast standard errors use the full covariance matrix of the two coefficients. All specifications include block, year-month, and municipality-by-year fixed effects plus property controls.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.28: Metro Proximity Threshold Sensitivity: Commercial Prices and Asking Rents

Threshold (m)	300	400	500	600	700	800
Panel A. Commercial sale prices						
Metro Only	0.011	0.057	0.059	0.076	0.015	−0.161
	(0.120)	(0.150)	(0.109)	(0.111)	(0.102)	(0.102)
Violence Only	−0.037	−0.036	−0.027	0.011	0.017	−0.014
	(0.058)	(0.061)	(0.064)	(0.066)	(0.068)	(0.071)
Both	0.166**	0.072	0.019	−0.041	−0.049	−0.079
	(0.081)	(0.065)	(0.062)	(0.063)	(0.061)	(0.063)
p: Violence Only = 0	0.525	0.551	0.673	0.868	0.804	0.844
p: Compounding	0.007	0.065	0.431	0.404	0.282	0.328
Panel B. Asking rents						
Metro Only	0.014	0.007	0.028	0.015	0.022	0.026*
	(0.013)	(0.012)	(0.018)	(0.016)	(0.014)	(0.015)
Violence Only	−0.015*	−0.015*	−0.012	−0.011	−0.015	−0.008
	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.010)
Both	−0.025	−0.018*	−0.021*	−0.027**	−0.012	−0.023**
	(0.016)	(0.010)	(0.011)	(0.011)	(0.011)	(0.010)
p: Violence Only = 0	0.083	0.081	0.160	0.236	0.112	0.386
p: Severity gradient	0.000	0.000	0.000	0.001	0.000	0.000
Observations	31,023 (Panel A); 674,479 (Panel B)					

Notes: Same specification as Table A.27 applied to commercial sale prices and asking rents. Commercial coefficients are imprecise throughout and show no consistent pattern across thresholds; the significant compounding contrast at the tightest cutoffs is not stable and is not interpreted. In the rental market the severity gradient between $V_b \geq 2$ and $V_b = 1$ blocks far from closed stations is significant at the 1% level at every threshold, but this reflects a positive coefficient for low exposure as much as a negative one for severe exposure (Section 5.3).
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A.29: Formal Pre-Shock Placebo: Fake Shock in April 2018

	Residential Prices	Asking Rents	Commercial Prices
Fake post × $V_b = 1$	0.0075	0.0006	−0.1010
	(0.0093)	(0.0077)	(0.0631)
Fake post × $V_b \geq 2$	−0.0010	−0.0051	0.0315
	(0.0115)	(0.0079)	(0.0532)
Observations	270,527	89,873	14,609
Pre / post split	131,487 / 139,040	32,366 / 57,507	7,819 / 6,790
Block FE	Yes	Yes	Yes
Year-month FE	Yes	Yes	Yes
Municipality × year FE	Yes	Yes	Yes

Notes: Each column re-estimates the baseline specification on pre-shock observations only, assigning a fake treatment date in April 2018, eighteen months before the actual shock. No coefficient approaches conventional significance. The uneven pre/post split in the rental column reflects growth in listing coverage over the sample period rather than a change in the underlying market. All specifications include property controls.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Figures

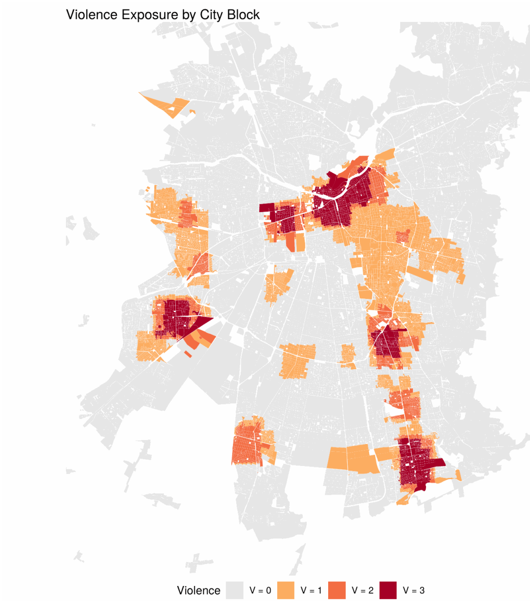

(a) Discrete violence index V_b

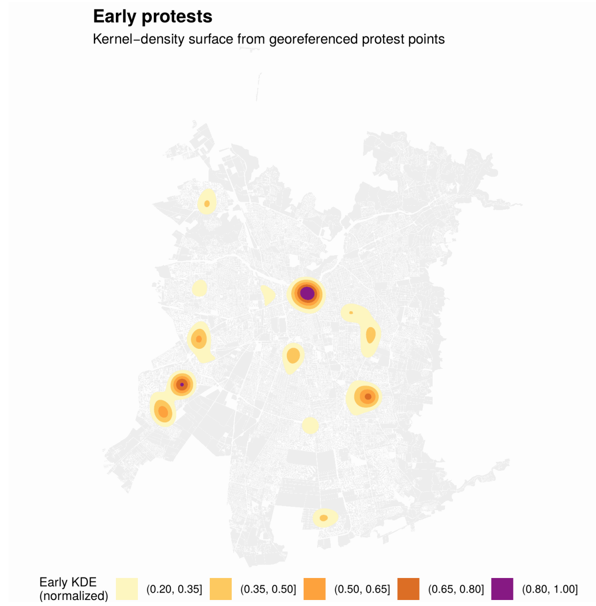

(b) Early protest geography, October 18–19, 2019

Figure A.1: Spatial anchoring of unrest at the outset. Panel (a) maps the final discrete violence index V_b based on Prosecutor hotspot intersections. Panel (b) maps a kernel-density surface constructed from georeferenced protest points observed on October 18–19, 2019. The comparison shows that the broad geography of final exposure was already visible during the first two days of the unrest. This early measure is used only as descriptive validation and not as the main treatment variable in the empirical analysis.

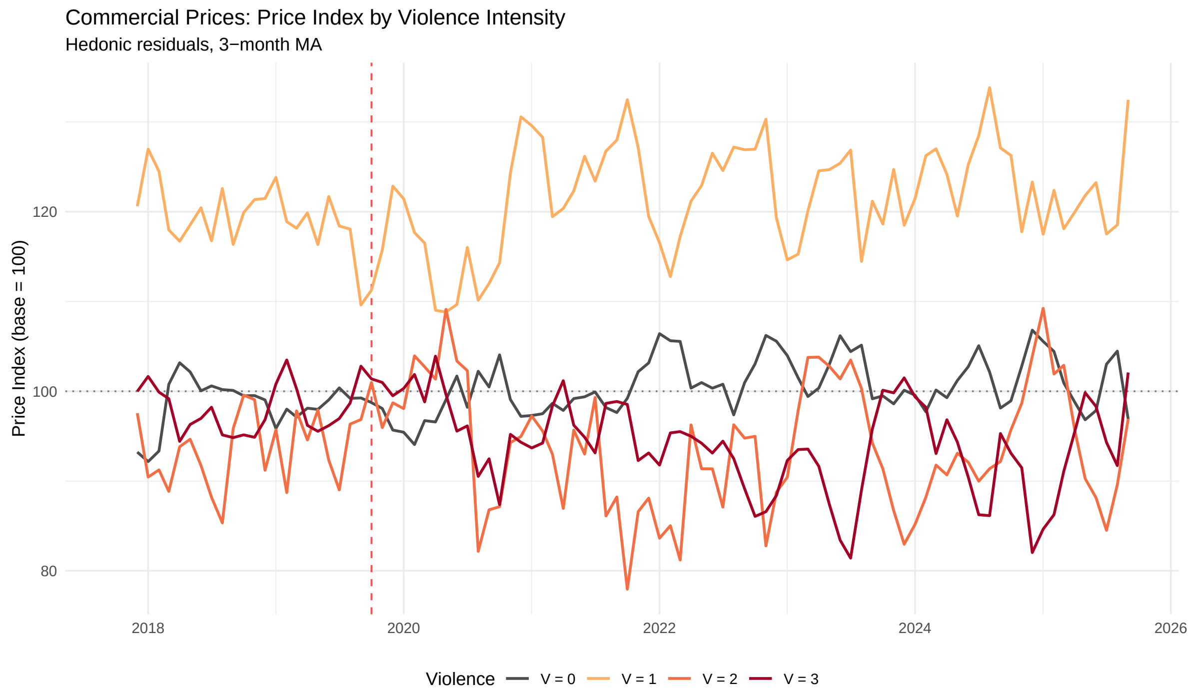

Figure A.2: Commercial Price Indices by Violence Intensity. Hedonic residuals from a regression of log price on property controls with block and year-month fixed effects, normalized to 100 in October 2017 and smoothed using 3-month moving averages. The vertical dashed line marks October 2019.

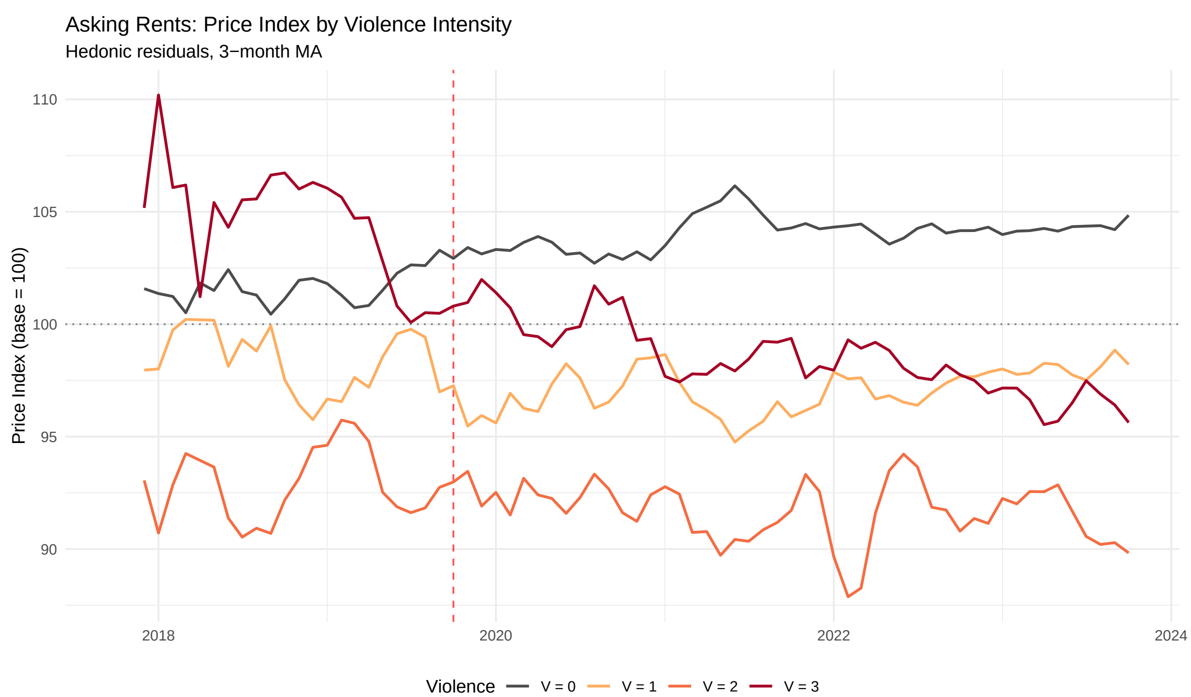

Figure A.3: Residential Rent Indices by Violence Intensity. Hedonic residuals from a regression of log price on property controls with block and year-month fixed effects, normalized to 100 in October 2017 and smoothed using 3-month moving averages. The vertical dashed line marks October 2019.

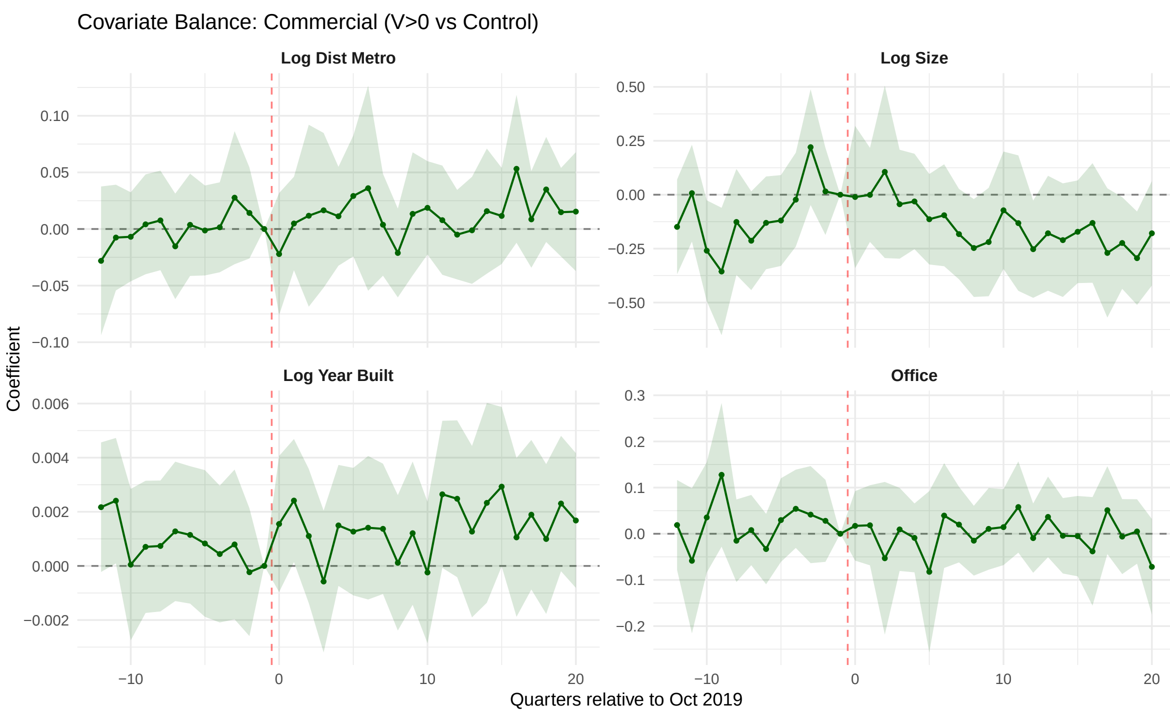

Figure A.4: Covariate Balance Event Studies (Commercial). Coefficients from regressions of each covariate on $\mathbb{1}\{V_b > 0\}\times$ quarter indicators with block and year-month fixed effects. Bands: 95% CIs, SEs clustered at block level. Reference: $Q{-}1$. Joint F-tests for post-shock coefficients: log size $p = 0.071$, log year built $p = 0.508$, apartment $p = 0.045$, log distance to metro $p = 0.425$.

Figure A.5: Covariate Balance Event Studies (Rents). Coefficients from regressions of each covariate on $\mathbb{1}\{V_b > 0\}\times$ quarter indicators with block and year-month fixed effects. Bands: 95% CIs, SEs clustered at block level. Reference: $Q{-}1$. Joint F-tests for post-shock coefficients: log size $p = 0.071$, log year built $p = 0.508$, apartment $p = 0.045$, log distance to metro $p = 0.425$.

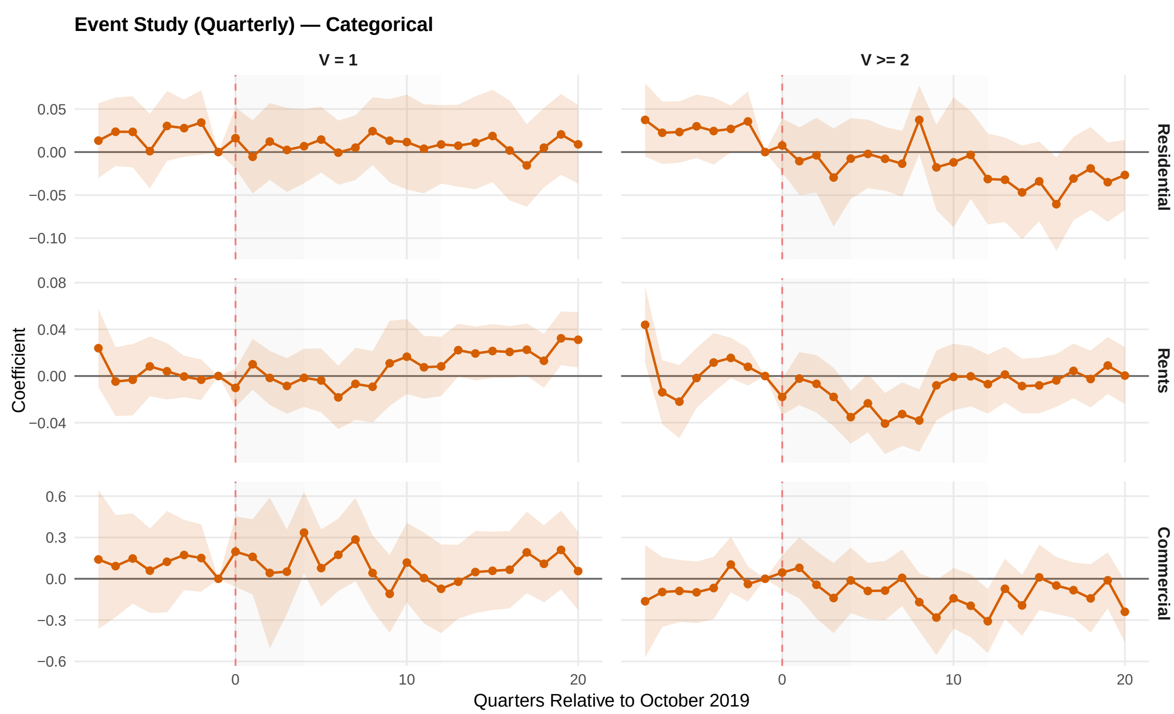

Figure A.6: Event Study (Quarterly): Prices and Rents. Same specification as Figure 6 at quarterly frequency. Pre-trend Wald p (residential): $V_b = 1$, $p = 0.513$; $V_b \geq 2$, $p = 0.442$.

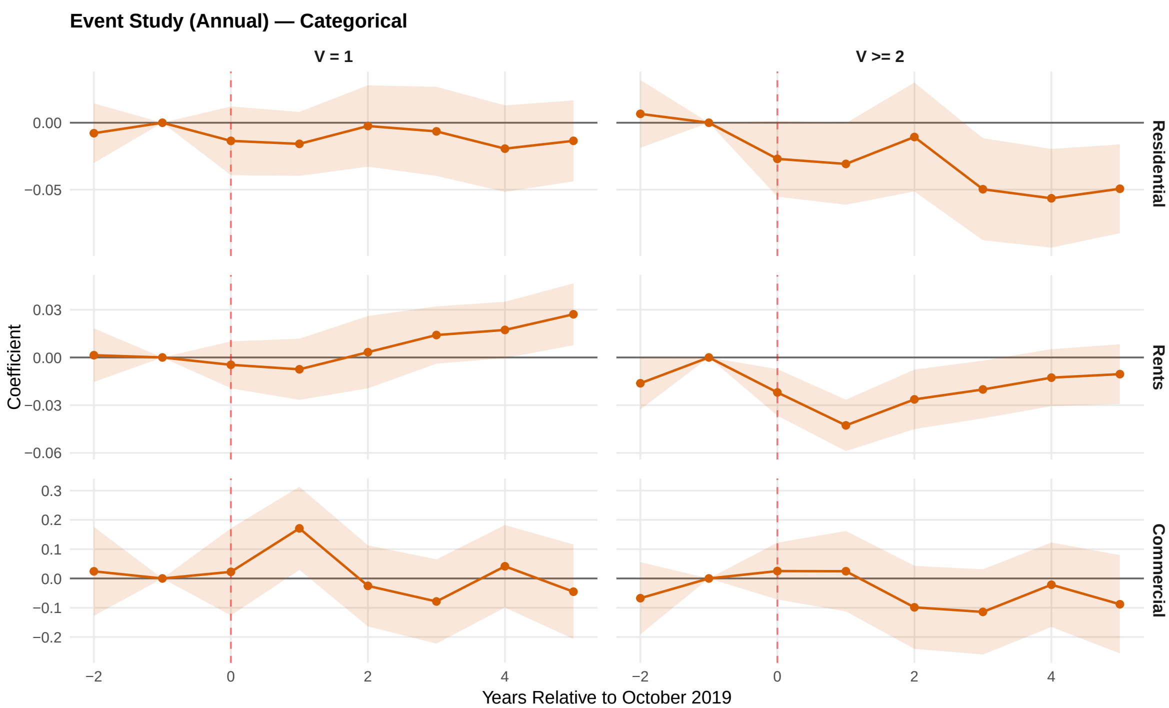

Figure A.7: Event Study (Annual): Prices and Rents. Same specification at annual frequency. Pre-trend Wald p (residential): $V_b = 1$, $p = 0.081$; $V_b \geq 2$, $p = 0.921$.