Mass Arrests and Local Public Life

Eleno Castro*

September 30, 2026

Abstract

El Salvador's 2022 mass-arrests crackdown provides rare evidence on how criminal governance shapes everyday life. I link national household survey data to pre-policy gang-related homicide exposure and compare post-crackdown changes across cantons. Previously exposed places report large gains in freedom of movement and ordinary local activity, closing much of the pre-policy gap with less exposed areas. Households also become more likely to use parks, sports courts, play areas, and community centers. The results reveal a hidden cost of criminal governance: restrictions on the community's ordinary use.

Keywords: Criminal governance; gangs; public security; mass arrests; mobility; public space; El Salvador.
JEL codes: K42, D74, O12, H56, R23.

*Department of Economics, Johns Hopkins University. Email: ecastr10@jhu.edu.

1 Introduction

In March 2022, following a weekend during which gangs killed 87 people (Associated Press, 2022), El Salvador declared a state of exception and initiated mass arrests of suspected gang members. Nearly 80,000 individuals were detained within two years (Alemán, 2024). The policy became one of Latin America's most prominent and controversial public-security interventions, combining large reported security gains with serious human-rights concerns (Human Rights Watch, 2025; Santos et al., 2026). It also provides a rare opportunity to evaluate one margin of this public-security intervention: whether households in previously gang-affected places regained effective access to local public goods and ordinary community life. This paper studies whether residents in previously gang-affected places regained the ability to conduct ordinary activities and use public and community spaces.

Public security matters for local welfare partly because it shapes whether households can safely use streets, parks, businesses, and community facilities. Criminal governance restricts this margin not only through victimization, migration, or income losses, but also by limiting ordinary community use. Existing work shows that gang-imposed borders restrict commuting and access to labor-market opportunities, while local crime and violence can reduce visits to businesses and neighborhoods (Melnikov et al., 2025; Fe and Sanfelice, 2022; Vachuska and Movahed, 2025). Much less is known about the reverse adjustment: whether residents recover the ordinary use of their communities when the state targets criminal organizations at scale. I study this margin using household reports of community freedom together with reported use of parks, sports courts, children's play areas, and community centers.

I find that communities previously exposed to gang-related violence experienced a substantial reopening of local public life. Before the crackdown, households in these places were less likely to report that residents could carry out ordinary activities in their communities, including going out at night, running a business, leaving the home unattended, allowing children to play outside, and women moving freely. After the crackdown, these gaps narrowed considerably across the different dimensions of reported community freedom. The pattern points to a broad-based improvement in ordinary community use rather than a change limited to a single survey item.

The observed change goes beyond reported freedom. Households in high-exposure cantons became 5.5 percentage points more likely to use at least one recreational or

community space. Measures of intensity, such as the number of spaces used and the usage rate among available spaces, also increased. Descriptive trends reveal a sharp rise following the onset of the crackdown, while event-study estimates indicate no clear differential improvement in high-exposure cantons immediately prior to 2022Q2. Because the crackdown was implemented nationwide, I interpret these estimates as reflecting differential gains in previously exposed areas rather than the policy's total effect.

The analysis integrates household records from the *Encuesta de Hogares de Propósitos Múltiples* (2019–2025) with canton-level exposure to gang-related violence prior to the crackdown. Exposure is defined as whether a canton recorded at least one gang-related homicide during 2016–2020. The empirical strategy compares post-2022Q2 changes in cantons with and without this pre-policy exposure, absorbing canton and year-quarter fixed effects. The estimated coefficient thus captures differential changes in areas where gang-related constraints on daily life were more likely to be binding before the policy.

I address four potential concerns regarding this interpretation. First, the results are robust when the pre-crackdown benchmark is anchored in pre-COVID survey years. Second, gang-related homicides are spatially concentrated in or near publicly mapped areas of gang control in San Salvador and Soyapango. Third, the community-freedom result remains positive when exposure is measured using homicide-rate intensity bins and when allowing urban and rural cantons to follow different post-crackdown paths. Finally, placebo outcomes and household characteristics do not exhibit patterns consistent with a generalized reporting or halo effect.

The gains are not uniform across households and localities. A systematic heterogeneity analysis finds that the community-freedom effect is larger among predicted high-impact profiles. These profiles are primarily characterized by household composition, including larger household size, more children and adolescents, a higher average number of male members aged 12–35, fewer older adults, and a lower prevalence of female-headed households. I interpret these findings as evidence of where gains are concentrated.

As a secondary check, I examine whether the reopening of local public life is reflected in selected labor-market outcomes. These results are exploratory. The estimates do not indicate broad increases in employment or labor-force participation. A limited number of work-related outcomes move in directions consistent with changes in mobility or

work organization, but these estimates are less central and less precise than the main public-life results. I therefore do not interpret them as evidence of a general expansion in labor-market activity.

The paper contributes to the public economics of security and local public goods by studying how a major security shock changed households' effective access to ordinary community life. The closest empirical evidence in this setting comes from Melnikov et al. (2025), who show that gang-controlled borders in San Salvador restrict commuting, freedom of movement, and access to labor-market opportunities. Related work shows that local crime and violence reduce consumer visits and neighborhood activity (Fe and Sanfelice, 2022; Vachuska and Movahed, 2025; Castro, 2026). Perceived safety can also affect consequential choices even without realized victimization: Borker (2021), for example, shows that women trade off college quality against safer travel in Delhi. Studies of crime and security in El Salvador examine the role of crime in international migration, the diffusion of gangs, and the crime effects of the 2022 crackdown (Abreha et al., 2025; Sviatschi, 2022; Santos et al., 2026). This paper studies a different question and direction of adjustment: whether large-scale enforcement against gangs reverses constraints on ordinary community life. It combines household reports of community freedom with use of parks, sports courts, play areas, and community centers—a within-community margin not observed in commuting, crime, migration, or digital foot-traffic outcomes. More broadly, the paper relates to work on territorial governance (Lessing, 2021; Blattman et al., 2025), gang organization and enforcement across institutional settings (Bruhn, 2021; Disney et al., 2026; Domínguez, 2026), extortion and local markets (Brown et al., 2025), and the effects of violence on firms and labor markets (Rozo, 2018; Velásquez, 2020; Utar, 2025).

The analysis does not estimate the crackdown's overall welfare effect or assess its legal, institutional, or human-rights consequences. Its contribution is narrower: it documents that large-scale public-security enforcement can rapidly alter a rarely measured margin of social and economic life—households' effective ability to move through, use, and participate in their local communities. This margin is economically meaningful because public security changes the set of local activities and public spaces that households can actually use.

2 Context, Data, and Measurement

The household data come from the *Encuesta de Hogares de Propósitos Múltiples* (EHPM), El Salvador's national household survey (Oficina Nacional de Estadística y Censos, 2026). The main repeated-cross-section analysis uses household records from 2019 through 2025. The survey reports the household's canton of residence, interview timing, survey weights, household characteristics, and a set of questions on local conditions and use of community facilities. I convert interview month to survey quarter for the descriptive and regression analyses.

Treatment is defined at the canton level using pre-crackdown gang-related homicide records obtained through public information requests to El Salvador's National Civil Police. A canton is classified as high exposure if it recorded at least one gang-related homicide during 2016–2020. These are the pre-crackdown years made available through the requests; later pre-policy records were requested but not provided. The longer pre-policy window also reduces sensitivity to isolated annual fluctuations in local violence. The main comparison is therefore between households living in cantons with and without recent pre-policy gang-related homicide exposure. In the main estimation sample, this definition classifies 418 of 746 cantons as high exposure, leaving 328 comparison cantons. I use "high-exposure cantons" as shorthand for this predetermined treatment group. The main estimation sample keeps cantons observed at least once before and once after the beginning of the crackdown in 2022Q2.

Because gang-related homicides are infrequent events, the five-year indicator may classify some gang-affected cantons as unexposed. Appendix Table A.5 and Figure A.5 assess this concern by splitting the homicide records into disjoint sets of years and using each resulting exposure indicator as an instrument for the other.

As a partial validation of the exposure measure, I compare gang-related homicides with publicly available maps of gang-controlled areas in San Salvador and Soyapango. Most gang-related homicides in these municipalities fall inside mapped gang-control polygons, and those outside the mapped polygons are on average located close to the nearest mapped gang-control area. This exercise is not a complete map of gang control nationwide, but it supports the interpretation that the homicide-based exposure measure captures areas connected to documented gang territorial control.

Because exposure is defined from pre-policy local conditions rather than random assignment, high- and low-exposure cantons differ in baseline characteristics. Appendix

Table A.1 reports pre-treatment weighted means by exposure status. High-exposure cantons report lower baseline community freedom and lower use of recreational and community spaces before the crackdown. These differences are important for interpretation: the empirical design does not compare outcome levels across exposure groups, but changes within cantons after 2022Q2 relative to contemporaneous changes in less exposed cantons.

The primary outcome is a community-freedom index. The index averages available household-level indicators for whether residents report that people in the community can go out at night, run a business, leave the home unattended, allow children to play outside, and that women can move freely, requiring at least three non-missing components. Higher values indicate greater reported freedom to use the community for ordinary activities. I also report estimates for each component to show whether the index is driven by one survey item or by a broader change across margins.

The second main set of outcomes measures use of recreational and community spaces. The EHPM records whether sports courts, parks or green areas, children's play areas, and community centers are available and, when available, whether the household uses them. The main indicators follow this questionnaire structure and measure use conditional on reported facility availability. From these items, I construct indicators for use of each type of space, an indicator for using any recreational or community space, the number of spaces used, and the use rate among available spaces. These outcomes complement the community-freedom index because they capture reported engagement with local spaces rather than only perceived freedom of movement. They are still survey-reported outcomes, but they are closer to realized behavior than the community-freedom questions.

As a robustness check, I also define public-space use without conditioning on reported facility availability. A household is coded as a nonuser when it reports that a particular facility is unavailable, while the outcome remains missing when the facilities module is not observed. Appendix Table A.4 reports these unconditional versions for each facility, any facility used, and the number of facility types used.

Sample sizes differ across outcomes because the EHPM records some items through different survey modules and skip patterns. In particular, the community-freedom index is based on broadly asked perceptions about local conditions, whereas recreational and community-space outcomes additionally depend on item-specific information about the availability and use of each facility type. I therefore estimate each outcome on the

sample with non-missing information for that measure.

Finally, I examine a limited set of labor-market outcomes as exploratory evidence on whether the reopening of local public life appears in productive margins. These outcomes capture labor supply, labor income among employed adults, workplace location, and job-search behavior. They are secondary to the main analysis. The paper does not interpret them as evidence of a broad labor-market response; instead, they are used to assess whether the public-life results are accompanied by suggestive changes in how work is organized.

3 Empirical Strategy

The empirical design compares changes after the onset of the crackdown across cantons with different pre-policy exposure to gang-related violence. Let Y_{hct} denote an outcome for household h in canton c and survey quarter t. Let $High_c$ be an indicator equal to one for cantons with at least one gang-related homicide in 2016–2020, and let $Post_t$ be an indicator equal to one from 2022Q2 onward. The baseline specification is

$$Y_{hct} = \beta(High_c \times Post_t) + \alpha_c + \lambda_t + \varepsilon_{hct}, \tag{1}$$

where α_c are canton fixed effects and λ_t are year-quarter fixed effects. The coefficient β measures the differential post-2022Q2 change in high-exposure cantons relative to cantons without pre-policy gang-related homicide exposure.

All household-level regressions use EHPM survey weights and cluster standard errors by canton, the level at which exposure status is assigned. The main repeated-cross-section sample is restricted to cantons observed at least once before and once after 2022Q2. Outcome-specific samples differ because some survey items are asked only for relevant subsamples or conditional on the availability of the corresponding facility. The identifying assumption is that, absent the crackdown, outcomes in high- and low-exposure cantons would have evolved in parallel conditional on canton and year-quarter fixed effects.

The design raises four concerns. First, the pre-period includes the COVID-19 pandemic, which directly affected mobility and the use of public spaces. I therefore report an auxiliary specification that uses pre-COVID survey years as the pre-crackdown reference period and excludes the COVID-period benchmark. Second, the exposure

measure is intentionally simple and may be correlated with canton size, urbanity, or the availability of public spaces. I address this concern with auxiliary checks that validate the homicide-based measure against public maps of gang-controlled areas, replace the binary treatment with homicide-rate intensity bins, and allow post-crackdown changes to vary with urban exposure. Third, the crackdown was national: comparison cantons were also exposed to the state of exception. The estimates should therefore be interpreted as differential gains in previously exposed areas, not as the total effect of the policy. Finally, because the outcomes are survey reports collected after a highly salient policy, differential reporting could mimic real changes. The household-composition outcomes in Appendix Table A.2 and the dwelling, service, and subsidy placebos in Appendix Table A.3 assess whether similarly reported outcomes shift in ways that would indicate a broader reporting or composition change.

To examine timing and pre-treatment dynamics, I estimate event-study specifications for the two main outcomes. Let $K(t)$ denote event time in survey years relative to 2022Q2. Event year 0 begins in 2022Q2, event year -1 is the omitted pre-crackdown year, and other event years are defined analogously. The event-study specification is

$$Y_{hct} = \sum_{k \neq -1} \beta_k \left(High_c \times 1\{K(t) = k\} \right) + \alpha_c + \lambda_t + \varepsilon_{hct}. \tag{2}$$

The coefficients β_k compare high- and low-exposure cantons in each event-year bin relative to the omitted pre-treatment period. Estimates for $k < 0$ provide a diagnostic for differential pre-trends, while estimates for $k \geq 0$ describe the evolution of post-crackdown differences. The event-study coefficients are grouped into event-year bins to improve precision and readability while reducing sensitivity to within-year seasonality, but the regressions continue to absorb year-quarter fixed effects.

The pre-COVID reference specification follows the same difference-in-differences structure as equation (1), but changes the reference period used to anchor pre-crackdown outcomes. Specifically, it excludes survey years 2020 and 2021 and 2022Q1, and adds 2017 and 2018 to the pre-period so that the omitted pre-treatment benchmark is based on pre-COVID survey years. This specification is useful because COVID directly affected the outcomes studied here; it tests whether the main results depend on comparing the post-crackdown period to a COVID-affected pre-period.

Finally, I summarize heterogeneity in the community-freedom effect using an honest GATES procedure. The procedure splits entire cantons into training and validation

samples, constructs a predicted-impact score using pre-policy or predetermined characteristics in the training sample, and estimates group-specific difference-in-differences effects in held-out validation cantons. The machine-learning step is used only to rank profiles out of sample; the group-specific estimates use the same difference-in-differences structure as equation (1). I interpret these results as evidence on where gains are concentrated, not as estimates of individual treatment effects or causal mechanisms for specific household characteristics. As an auxiliary validation reported in the appendix, I transport the predicted-impact score to a smaller panel of reobserved households to assess whether the score ordering is directionally preserved outside the repeated cross-section.

4 Effects on Local Public Life

The main results point to a substantial reopening of local public life in cantons previously exposed to gang-related violence. High-exposure cantons entered the period with lower reported community freedom and lower use of recreational and community spaces than less exposed cantons. The estimates below show that these gaps narrowed sharply after the onset of the crackdown.

Table 1 reports the baseline difference-in-differences estimates. The community-freedom index rises by 0.077 points after 2022Q2 in high-exposure cantons relative to other cantons. Since the index is bounded between zero and one and is constructed from up to five binary indicators, this is naturally interpreted as a 7.7-point increase in reported community freedom. Relative to the pre-treatment mean in exposed cantons, the estimate represents an increase of 10.5 percent. The magnitude is large relative to the pre-policy gap between exposure groups, indicating that the crackdown closed much of the initial difference in reported access to local public life.

The component estimates show that this result is broad rather than driven by a single survey item. Households in previously exposed cantons become more likely to report that residents can go out at night, run a business, leave the home unattended, allow children to play outside, and that women can move freely. The estimated gains across these components are all positive and statistically significant, ranging from roughly six to ten percentage points. This breadth is important for interpretation: the change appears across several margins of everyday public life rather than only in one dimension of perceived security.

Panel B of Table 1 shows that these reported freedoms are accompanied by greater use of local spaces. Households in high-exposure cantons become 5.5 percentage points more likely to report using at least one recreational or community space, a 10.4 percent increase relative to the pre-treatment mean in exposed cantons. The number of spaces used and the use rate among available spaces also increase. Among individual facilities, the estimates are positive for sports courts, parks or green areas, children's play areas, and community centers; the first three are statistically significant, while the community-center estimate is positive but less precise. I therefore interpret the aggregate measures as the main evidence on public-space use, with the individual facility estimates showing that the pattern is not concentrated in only one type of space.

Appendix Table A.4 shows that this result is not driven by conditioning on reported facility availability. When households reporting an unavailable facility are retained as nonusers, the effect on use of any recreational or community space is 5.0 percentage points (SE 1.4 $N = 118,148$), close to the conditional estimate of 5.5 percentage points. The number of facility types used rises by 0.117. The community-center estimate is smaller but more precise under the unconditional definition: 2.2 percentage points (SE 1.1) rather than 3.3 percentage points (SE 2.1) in the conditional specification.

Figure 1 examines the timing of these changes. Panels A and B plot descriptive quarterly series for the two main outcomes. These descriptive panels do not identify causal effects, but they show that the aggregate timing of the change lines up closely with the policy shock: both community freedom and recreational/community space use rise after the beginning of the crackdown.

Panels C and D show the corresponding event-study estimates. For the community-freedom index, there is no clear evidence that high-exposure cantons were already improving relative to comparison cantons before the crackdown. After 2022Q2, the coefficients turn positive and remain positive in later event years. The event-study pattern for recreational and community space use is less precise, but it follows the same broad timing: the pre-event coefficients do not show differential improvement in high-exposure cantons, while the post-event estimates become positive, especially in later event years. The event-study evidence therefore supports the interpretation that the main estimates capture a post-crackdown change rather than a continuation of a pre-existing improvement in high-exposure areas.

Having established the baseline pattern, I next examine whether it is sensitive to the

pre-period benchmark, exposure construction, urbanity, or reporting-related concerns. A natural concern is that the baseline pre-period includes the COVID-19 pandemic, which directly affected mobility and the use of public spaces. Table 2 addresses this concern by comparing the baseline estimates with an auxiliary specification that anchors the pre-crackdown benchmark in pre-COVID survey years. The estimates are not smaller under this alternative benchmark. The community-freedom effect rises to 0.087, and the effect on any recreational or community space use rises to 0.074. This reduces the concern that the main results are an artifact of comparing the post-crackdown period only to a COVID-affected pre-period.

Because the binary exposure indicator may be noisy, Appendix Table A.5 implements a split-year IV sensitivity check. Across 500 random within-canton splits, the median estimates are 0.154 for community freedom and 0.117 for use of any recreational or community space, compared with baseline estimates of 0.077 and 0.055; the minimum first-stage F-statistics are 41.4 and 41.0. The estimates remain positive across splits and are consistently larger than the baseline estimates. The same pattern appears in all ten common-calendar partitions and both instrument directions (Appendix Figure A.5). I interpret this exercise as evidence consistent with attenuation from exposure misclassification, rather than as a definitive correction for latent gang control.

Additional appendix checks support the main interpretation while clarifying the scope of the results. Replacing the binary exposure indicator with homicide-rate intensity bins yields positive effects on community freedom across the exposure distribution, although the pattern is not monotonic. This is consistent with the fact that homicide intensity is an imperfect proxy for territorial control: high homicide rates may reflect contested violence as well as stable gang control. The recreational-space margin is more sensitive and appears more concentrated in urban settings, which is consistent with the availability and salience of the public facilities measured in the survey. To assess whether the main results could reflect changes in survey composition or a broad pro-government or optimism-related reporting shift, I examine respondent-composition and placebo outcomes. The respondent-composition coefficients in Appendix Table A.2 are jointly insignificant with survey weights ($p = 0.388$) and without them ($p = 0.171$). The dwelling, service, and subsidy placebos in Appendix Table A.3 likewise do not show broad shifts comparable in size and timing to the main public-life effects.

Overall, the evidence points to a reopening of local public life in previously exposed cantons. Because the crackdown was national, the estimates should not be interpreted

as showing that less exposed cantons were unaffected. Instead, they show that previously exposed cantons experienced larger gains, closing much of the pre-policy gap in reported community freedom visible in Appendix Figure A.6, and increasing engagement with local public and community spaces. The paper therefore identifies a specific effect of the security shock: a substantial change in households' reported ability to move through, use, and participate in their local communities.

5 Heterogeneity and Exploratory Economic Margins

The average effect on community freedom masks meaningful heterogeneity across household and locality profiles. Figure 2 summarizes this heterogeneity using an honest GATES procedure. The procedure is used only to rank profiles by predicted impact in held-out cantons; the group-specific effects are then estimated with the same difference-in-differences structure as the main analysis. Details on the score construction, sample splitting, and profile comparisons are reported in Appendix Section A.2.

Two caveats apply. The procedure adapts methods developed for randomized experiments to a difference-in-differences setting, so each group-specific estimate inherits the parallel-trends assumption of the baseline design. And because the data are repeated cross-sections, profile characteristics describe who lives in high-gain groups, not mechanisms: I interpret the results as evidence on where gains are concentrated, not on why.

Panel A shows an increasing pattern across predicted-impact groups. The lowest predicted-impact group exhibits a smaller estimated effect on the community-freedom index, while the point estimates rise across groups and are largest for the highest predicted-impact group. The split-stability intervals show a similar ordering across honest sample splits. This pattern suggests that the average gain in community freedom is not uniform; it is concentrated among profiles predicted ex ante to experience larger improvements in local public life.

Panel B describes the high and low predicted-impact profiles. The largest profile differences are in household composition. Relative to the low predicted-impact group, the high predicted-impact group is less likely to be female-headed, has larger households, contains more children and adolescents, has a higher average number of male members ages 12–35, and has fewer older adults. The dependency ratio is also higher in the

high predicted-impact group. These characteristics describe households for whom changes in local public life may be especially salient: children and adolescents can use outdoor and community spaces, young adults may face local mobility constraints, and larger households may have more members whose daily routines depend on the neighborhood environment.

The CLAN differences should not be interpreted as causal mechanisms. Household-composition variables in a repeated cross-section may partly reflect changes in the composition of households observed in each group rather than fixed household attributes. The heterogeneity exercise therefore provides a disciplined summary of which profiles experience larger gains, not a decomposition of the channels through which those gains occur.

As an auxiliary validation, Appendix Figure A.1 asks whether the ranking learned in the repeated cross-section also has content in a different sample structure. I transport the predicted-impact score to a smaller panel of reobserved households and estimate within-household changes by score group. The point estimates are ordered in the same direction, with larger gains among higher-score households, although the confidence intervals are wide. I therefore interpret this exercise as a directional check that the score is not only capturing contemporaneous composition in the repeated cross-section.

The labor-market results provide a separate, more exploratory margin. If greater community freedom relaxed constraints on local movement, one might expect some changes in work-related behavior even in the absence of broad employment effects. Appendix Table A.9 examines this possibility for selected person-level outcomes. The estimates do not show broad increases in employment or labor-force participation. This is important for interpretation: the main results should not be read as evidence of a general labor-market expansion.

The few margins with suggestive movement are closer to mobility and work organization than to extensive-margin labor supply. The point estimates for workplace-location outcomes move in a direction consistent with changes in where or how existing work is organized. I therefore interpret these estimates as suggestive evidence that improved local freedom may have affected some work-related margins, not as evidence of broad labor-market gains.

6 Conclusion

This paper studies whether a large public-security shock changed everyday access to local public life. Using household survey data from El Salvador and pre-policy variation in canton-level exposure to gang-related violence, I find that the 2022 mass-arrests crackdown was followed by large differential gains in reported community freedom in previously exposed cantons. These cantons began the period with substantially lower reported access to local public life; after the crackdown, they closed much of that pre-policy gap. The same places also report greater use of recreational and community spaces.

The results point to a specific margin of adjustment: a reopening of ordinary community life. The strongest evidence is on community freedom, where the estimates are large, broad across components, and aligned with the timing of the crackdown. The increase in recreational and community space use provides a complementary behavioral margin. The public-space estimates remain similar when households reporting an unavailable facility are retained as nonusers. The results are also robust to using pre-COVID years as the reference period, reducing the concern that the main estimates simply reflect recovery from a pandemic-period benchmark.

Because the crackdown was national, the estimates should not be interpreted as the total effect of the policy. Less exposed cantons also experienced the state of exception and may have improved after 2022Q2. The estimates instead capture the additional gains in places where pre-policy gang-related constraints on local public life were likely more salient. This interpretation is supported by auxiliary checks showing that the homicide-based exposure measure is related to publicly mapped areas of gang control, and that the main community-freedom result is not driven by broad changes in household composition or reporting behavior.

The evidence is more limited on broader economic margins. The heterogeneity analysis suggests that gains are concentrated among household and locality profiles for which local public life may be especially salient, but these profiles should not be interpreted as causal mechanisms. Exploratory labor-market results do not show broad increases in employment or labor-force participation. At most, selected work-related margins move in directions consistent with changes in mobility or work organization. These results are suggestive and secondary to the main finding.

The paper does not estimate the overall welfare effect of the crackdown, nor does

it evaluate its legal, institutional, or human-rights consequences. Its contribution is narrower: it documents that a large security shock can rapidly change households' reported ability to move through, use, and participate in their local communities. This margin matters because the costs of criminal governance are not limited to victimization, migration, or income losses. They also include restrictions on the ordinary use of the community itself.

7 Tables and Figures

Table 1: Effects on Community Freedom and Use of Local Public Spaces

Panel A. Community freedom index and components							
Outcome	β	SE	Mean	%Δ	N	Clusters	R^2
Can go out at night	0.091***	0.014	0.696	13.1	115,509	746	0.148
Can run a business	0.084***	0.016	0.706	11.9	115,509	746	0.138
Can leave the house alone	0.059***	0.013	0.752	7.9	115,509	746	0.103
Children can play outside	0.098***	0.017	0.709	13.8	62,487	746	0.126
Women can move freely	0.065***	0.012	0.795	8.2	105,835	746	0.106
Community freedom index	0.077***	0.013	0.733	10.5	115,509	746	0.164
Panel B. Use of recreational and community spaces							
Outcome	β	SE	Mean	%Δ	N	Clusters	R^2
Sports court used	0.042***	0.012	0.459	9.1	72,200	746	0.053
Park or green area used	0.083***	0.020	0.495	16.9	29,095	637	0.071
Children's play area used	0.060**	0.023	0.380	15.7	22,237	538	0.056
Community center used	0.033	0.021	0.467	7.0	30,483	678	0.115
Any recreational space used	0.055***	0.013	0.526	10.4	78,279	746	0.068
Number of recreational spaces used	0.147***	0.036	0.962	15.3	78,279	746	0.109
Use rate among available recreational spaces	0.048***	0.011	0.452	10.5	78,279	746	0.061

Notes: Panel A reports difference-in-differences estimates for the community-freedom index and its components. Panel B reports estimates for use of recreational and community spaces. The coefficient is the interaction between the post-2022Q2 period and pre-policy gang-related homicide exposure. Mean is the pre-treatment mean among exposed cantons. Standard errors are clustered by canton.

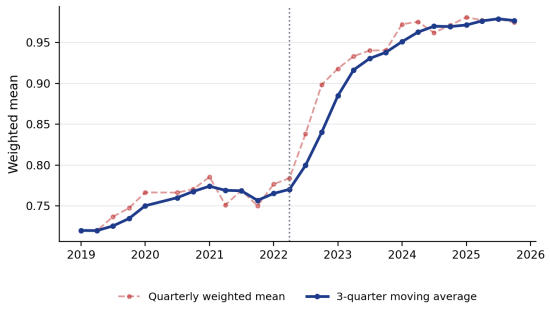
Table 2: Repeated Cross-Section Estimates

	Baseline RCS		Pre-COVID reference	
	Freedom	Any Rec. Used	Freedom	Any Rec. Used
High exposure × Post	0.077***	0.055***	0.087***	0.074***
	(0.013)	(0.013)	(0.014)	(0.014)
Mean	0.733	0.526	0.676	0.530
Change %	10.5	10.4	12.8	14.0
Clusters	746	746	687	687
Obs.	115,509	78,279	117,467	81,575
R^2	0.164	0.068	0.205	0.063

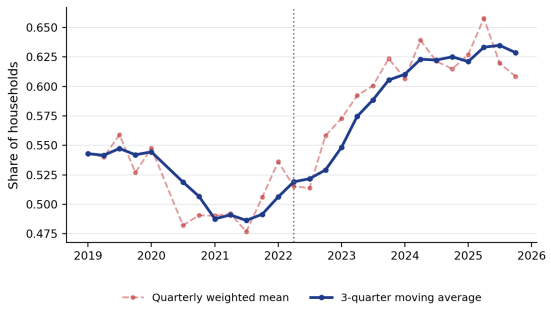
Notes: The table compares the baseline repeated cross-section DID estimates with an auxiliary repeated cross-section specification that uses only pre-COVID years to anchor the pre-crackdown reference period. The auxiliary specification excludes survey years 2020 and 2021 and 2022Q1, and adds 2017 and 2018 to the pre-period so that the omitted pre-treatment benchmark is based on pre-COVID survey years. Repeated cross-section estimates follow cantons over time. Mean is the pre-treatment mean among exposed cantons. Standard errors are clustered by canton.

Figure 1: Descriptive Trends and Event-Study Estimates for Local Public Life

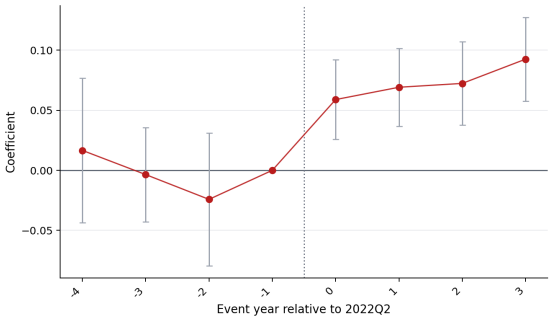
Panel A. Community freedom index, descriptive series



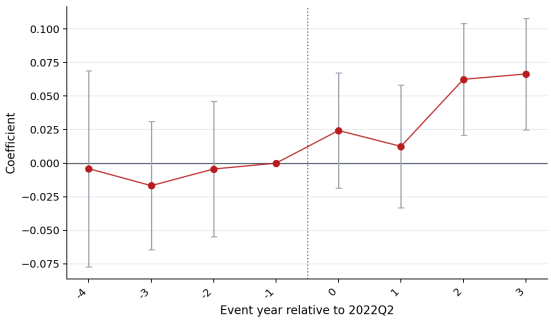
Panel B. Recreational/community space use, descriptive series



Panel C. Community freedom index, event study

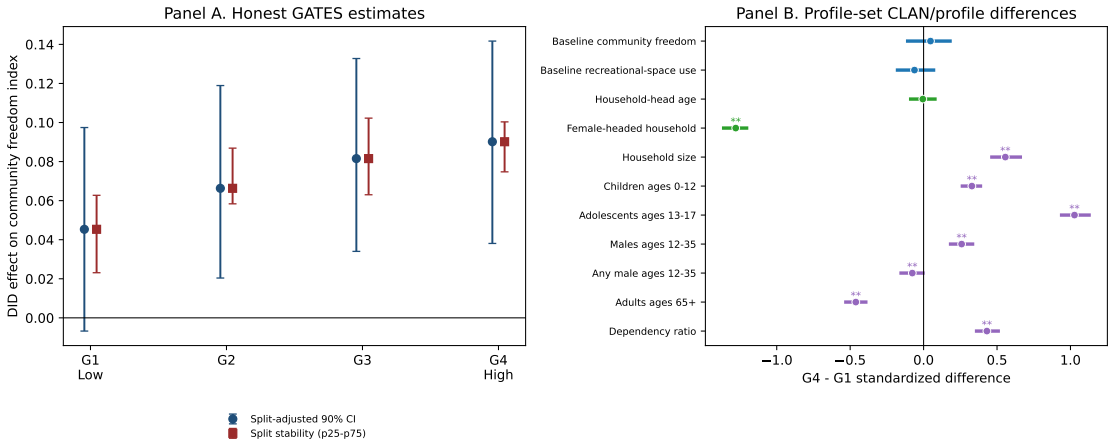


Panel D. Any recreational/community space used, event study



Notes: Panels A and B report descriptive quarterly series for the two main local public-life outcomes. Panels C and D report event-study estimates for the same outcomes, comparing high-exposure cantons with other cantons around the onset of the 2022 crackdown.

Figure 2: Heterogeneous Effects on Community Freedom



Notes: Panel A reports grouped average treatment-effect estimates for community freedom by predicted-impact group. Blue circles show split-adjusted confidence intervals constructed from split-specific clustered confidence intervals. Red squares show the interquartile range of split-specific point estimates and summarize stability across honest sample splits; they are not confidence intervals. Panel B summarizes baseline characteristics of households and localities across predicted-impact groups using split-adjusted confidence intervals.

Declaration of Generative AI and AI-Assisted Technologies in the Manuscript Preparation Process

During the preparation of this work, the author used OpenAI's ChatGPT to assist with manuscript drafting, language editing, and organization of arguments. The tool was also used for limited coding support, such as code structuring and debugging. After using this tool, the author reviewed, edited, and verified the content as needed. All empirical analysis, data processing, results, and interpretations were independently checked and finalized by the author. No restricted administrative microdata were provided to the AI tool. The author takes full responsibility for the content of the published article.

References

Abreha, K., Johnson, T., and Robertson, R. (2025). Is crime a "root cause" of Central American emigration? evidence from El Salvador. *Journal of Development Economics*, 174:103456.

Alemán, M. (2024). El Salvador extends anti-gang emergency decree for 24th time. it's now been in effect for two years. *Associated Press*. Available at https://apnews.com/article/el-salvador-gang-crackdown-emergency-decree-154d50d40d1f2a46a48b6880624df141. Accessed May 2026.

Associated Press (2022). El Salvador locks down prisons after wave of 87 killings over weekend. *The Guardian*. Available at https://www.theguardian.com/world/2022/mar/28/el-salvador-prisons-gang-killings. Accessed May 2026.

Athey, S. and Imbens, G. (2016). Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences*, 113(27):7353–7360.

Blattman, C., Duncan, G., Lessing, B., and Tobón, S. (2025). Gang rule: Understanding and countering criminal governance. *The Review of Economic Studies*, 92(3):1497–1531.

Borker, G. (2021). Safety first: Perceived risk of street harassment and educational choices of women. Policy Research Working Paper 9731, World Bank, Washington, DC.

Brown, Z. Y., Montero, E., Schmidt-Padilla, C., and Sviatschi, M. M. (2025). Market structure and extortion: Evidence from 50,000 extortion payments. *The Review of Economic Studies*, 92(3):1595–1624.

Bruhn, J. (2021). Competition in the black market: Estimating the causal effect of gangs in chicago. Orlando Bravo Working Paper 2021-004, Brown University, Department of Economics.

Castro, E. (2026). Violence reduction, firms, and local labor markets: Evidence from a ceasefire. Working paper. Available at https://www.elenocastro.com/assets/pdf/CeasefireAccessed June 2026.

Chernozhukov, V., Demirer, M., Duflo, E., and Fernández-Val, I. (2025). Fisher–schultz lecture: Generic machine learning inference on heterogeneous treatment effects in randomized experiments, with an application to immunization in India. *Econometrica*, 93(4):1121–1164.

Disney, R., Kirchmaier, T., Machin, S., and Villa, C. (2026). Gangs of london and public housing. CEP Discussion Paper 2147, Centre for Economic Performance, London School of Economics.

Domínguez, M. (2026). Sweeping up gangs: The effects of tough-on-crime policies from a network approach. *American Economic Journal: Economic Policy*, 18(1):127–158.

Fe, H. and Sanfelice, V. (2022). How bad is crime for business? evidence from consumer behavior. *Journal of Urban Economics*, 129:103448.

Fundación Padre Arrupe (2021). El Salvador. Map caption attributes source to El Salvador's Policía Nacional Civil. Available at https://www.fundacionpadrearrupe.com/el-salvador. Accessed May 2026.

Human Rights Watch (2025). World report 2025: El Salvador. Available at https://www.hrw.org/world-report/2025/country-chapters/el-salvador. Accessed May 2026.

InSight Crime (2019). Extorsión y presencia de pandillas en San Salvador. Available at https://insightcrime.org/es/investigaciones/extorsion-presencia-de-pandillas-en-san-salvador/. Accessed May 2026.

Lessing, B. (2021). Conceptualizing criminal governance. *Perspectives on Politics*, 19(3):854–873.

Melnikov, N., Schmidt-Padilla, C., and Sviatschi, M. M. (2025). Gangs, labor mobility, and development. *Econometrica*, 93(6):2083–2121.

Oficina Nacional de Estadística y Censos (2026). Encuesta de hogares de propósitos múltiples (EHPM). Available at https://onec.bcr.gob.sv/encuesta-de-hogares-de-propositos-multiples-ehpm/. Accessed June 2026.

Rozo, S. V. (2018). Is murder bad for business? evidence from Colombia. *The Review of Economics and Statistics*, 100(5):769–782.

Santos, M., Silverio-Murillo, A., Balmori-de-la-Miyar, J. R., and Rodríguez, A. (2026). Crime impacts of El Salvador's crackdown policy. Working Paper 20, School of Government and Public Transformation, Tecnológico de Monterrey. Available at https://ssrn.com/abstract=6116128. Accessed June 2026.

Sviatschi, M. M. (2022). Spreading gangs: Exporting US criminal capital to El Salvador. *American Economic Review*, 112(6):1985–2024.

Utar, H. (2025). Firms and labor in times of violence: Evidence from the Mexican drug war. *The World Bank Economic Review*, 39(3):632–662.

Vachuska, K. and Movahed, M. (2025). The causal effect of gun violence on everyday mobility patterns across US neighborhoods. *Spatial Demography*, 13(1):7.

Velásquez, A. (2020). The economic burden of crime: Evidence from Mexico. *Journal of Human Resources*, 55(4):1287–1318.

Wager, S. and Athey, S. (2018). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242.

A Additional Tables and Figures

A.1 Pre-Treatment Differences Across Exposure Zones

Table A.1: Pre-Treatment Differences Across Exposure Zones

Variable	Low exposure	High exposure
Household composition		
Household size	3.394	3.344
	(0.018)	(0.012)
Children ages 0-12	0.735	0.654
	(0.009)	(0.007)
Adolescents ages 13-17	0.293	0.273
	(0.006)	(0.004)
Males ages 12-35	0.686	0.646
	(0.008)	(0.006)
Any male ages 12-35	0.531	0.501
	(0.005)	(0.004)
Members ages 65+	0.335	0.339
	(0.006)	(0.005)
Dependency ratio	0.588	0.547
	(0.007)	(0.005)
Household-head profile		
Household-head age	49.658	50.525
	(0.179)	(0.125)
Female-headed household	0.357	0.388
	(0.005)	(0.004)
Locality baseline profile		
Baseline community freedom	0.807	0.722
	(0.001)	(0.001)
Baseline recreational-space use	0.557	0.483
	(0.002)	(0.001)

Notes: The table reports weighted pre-treatment means for households in cantons with and without at least one gang-related homicide during 2016–2020. Standard errors of the mean are reported in parentheses. The table is descriptive and is intended to characterize the exposure groups before the crackdown.

A.2 Details on Heterogeneity Procedure

This section describes the machine-learning-assisted heterogeneity exercise used in Figure 2. The goal is to summarize whether the average effect on the community-freedom index is concentrated among particular household and locality profiles. The exercise follows the logic of honest heterogeneous-treatment-effect methods, in which one part of the data is used to learn a prediction rule and held-out observations are used to estimate effects for the resulting groups (Athey and Imbens, 2016; Wager and

Athey, 2018). It also follows the GATES and CLAN framework of Chernozhukov et al. (2025), using predicted-impact groups to summarize treatment-effect heterogeneity and profile comparisons to describe the most and least affected groups. The exercise does not redefine treatment status and does not replace the baseline difference-in-differences design. The identifying variation remains the post-2022Q2 differential change between cantons with and without pre-policy gang-related homicide exposure.

The main heterogeneity exercise uses a profile set of pre-policy or predetermined characteristics. The score includes the baseline components of the community-freedom index, baseline recreational-space use, household-head age, female-headed household status, and household-composition variables. Treatment indicators, post-period indicators, treatment-by-post interactions, identifiers, time variables, the outcome, and post-treatment outcomes or mechanisms are excluded from score construction.

The procedure is honest at the canton level. In each split, entire cantons are assigned either to a training sample or to a validation sample. This preserves the level at which exposure is assigned and prevents observations from the same canton from being used both to learn the heterogeneity ranking and to estimate group-specific effects. The training sample is used only to construct the predicted-impact score. The validation sample is then used to sort households into predicted-impact groups and estimate group-specific difference-in-differences effects. This separation between ranking and estimation is the source of the "honesty" in the procedure.

Let Y_{hct} denote the community-freedom index for household h in canton c and survey quarter t. Let $High_c$ be an indicator for cantons with at least one gang-related homicide in 2016–2020, and let $Post_t$ be an indicator for survey quarters from 2022Q2 onward. Let X_{hct} denote the profile variables used to rank households. In the training sample, the procedure estimates four conditional mean functions:

$$m_{11}(x) = E[Y_{hct} \mid High_c = 1, Post_t = 1, X_{hct} = x], \tag{3}$$

$$m_{10}(x) = E[Y_{hct} \mid High_c = 1, Post_t = 0, X_{hct} = x], \tag{4}$$

$$m_{01}(x) = E[Y_{hct} \mid High_c = 0, Post_t = 1, X_{hct} = x], \tag{5}$$

$$m_{00}(x) = E[Y_{hct} \mid High_c = 0, Post_t = 0, X_{hct} = x]. \tag{6}$$

The fitted functions define a predicted difference-in-differences score:

$$\widehat{S}_{hct} = [\widehat{m}_{11}(X_{hct}) - \widehat{m}_{10}(X_{hct})] - [\widehat{m}_{01}(X_{hct}) - \widehat{m}_{00}(X_{hct})]\,. \tag{7}$$

This score is used only as a ranking device. It is not interpreted as an individual treatment effect.

The fitted score is applied to observations in the validation sample. Households in held-out cantons are sorted into predicted-impact groups, $G_1, \ldots, G_Q$, using the weighted pre-treatment distribution of the score. The lowest-score group is denoted G_1 and the highest-score group is denoted G_Q. Cutoffs are computed using the validation-sample pre-treatment distribution and then applied to both pre- and post-treatment validation observations. This keeps group definitions from being mechanically determined by post-treatment changes in the outcome.

For each predicted-impact group, I estimate a group-specific difference-in-differences effect in the validation sample. The implementation is equivalent to estimating

$$Y_{hct} = \sum_{q=1}^{Q} \beta_q \left(High_c \times Post_t \times 1\{G_{hct} = q\} \right) + \alpha_{cq} + \lambda_{tq} + \varepsilon_{hct}, \tag{8}$$

where α_{cq} are group-specific canton fixed effects and λ_{tq} are group-specific year-quarter fixed effects. Equivalently, the model can be estimated as separate difference-in-differences regressions within each predicted-impact group. The coefficients β_q are grouped average treatment-effect estimates sorted by predicted impact. The main contrast compares the highest predicted-impact group with the lowest predicted-impact group.

The procedure is repeated across multiple canton-level splits. Figure 2 reports the grouped estimates across splits. The point estimates summarize the central tendency of the split-specific group effects. The figure reports both split-adjusted confidence intervals, constructed from split-specific canton-clustered confidence intervals, and interquartile ranges of the split-specific point estimates. This approach shows both statistical uncertainty within splits and stability across honest splits. The split-stability intervals have a different interpretation from confidence intervals. They report the interquartile range of the point estimates obtained across honest canton-level splits. They are intended to show whether the ranking of predicted-impact groups is stable to the particular training-validation partition of cantons. They should not be interpreted as statistical confidence intervals. The confidence intervals instead summarize sampling uncertainty within split-specific clustered regressions, while the stability intervals summarize sensitivity to the random split.

Panel B of Figure 2 reports a classification analysis, following the CLAN logic in Chernozhukov et al. (2025), that describes the high and low predicted-impact groups. For each characteristic X_j, the comparison is

$$\Delta_j^{CLAN} = E_w[X_j \mid G = Q, Post = 0] - E_w[X_j \mid G = 1, Post = 0], \tag{9}$$

where $E_w[\cdot]$ denotes a survey-weighted mean. The comparison is made in the held-out validation sample and, when applicable, is restricted to pre-treatment observations. These differences describe the profiles associated with larger predicted gains; they are not causal effects of the corresponding characteristics.

The heterogeneity exercise inherits the identifying assumptions of the baseline difference-in-differences design. In particular, for each predicted-impact group, the untreated change in community freedom would need to have evolved similarly in high- and low-exposure cantons absent the post-2022Q2 security shock. The machine-learning step provides an honest and reproducible way to define groups using pre-policy information, but it does not by itself solve violations of parallel trends. For this reason, the heterogeneity results are interpreted as evidence on where the average effect is concentrated, not as estimates of household-level treatment effects.

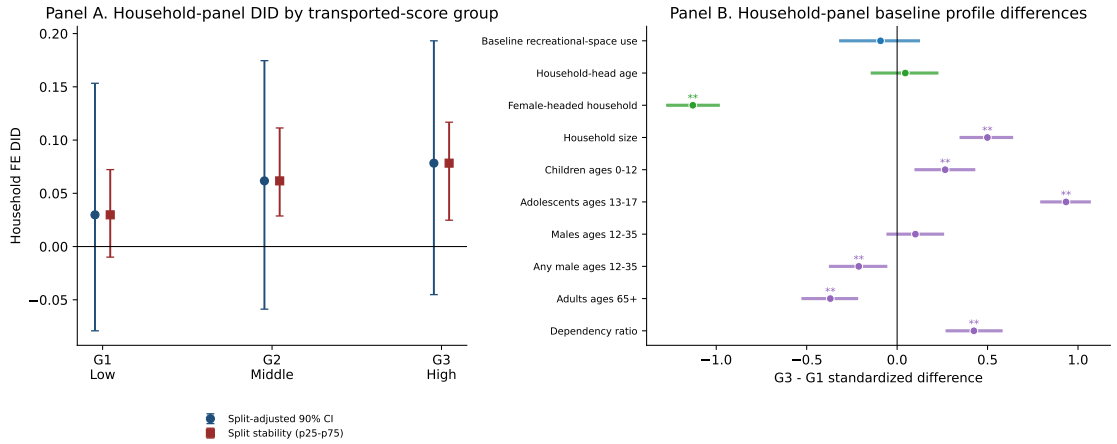
Household-composition variables require additional caution. In a repeated cross-section, differences in household composition may reflect changes in which households are observed in each group, not only fixed characteristics of the same households. I therefore interpret variables such as children, the average number of male members ages 12–35, older adults, and household size as descriptive profile characteristics. They help characterize the high predicted-impact group but do not establish mechanisms.

As an additional validation, I transport the predicted-impact score to a panel of reobserved households. This exercise asks whether a score learned in the repeated cross-section can also order households by within-household gains in community freedom. For each split, cantons used for household-panel validation are excluded from score training. The score is trained in the repeated cross-section using baseline household and locality covariates, and the fitted score is then applied to reobserved households in the validation cantons. Households are sorted into predicted-impact groups using their baseline covariates, and household fixed-effect difference-in-differences estimates are computed by score group.

Figure A.1 reports the panel validation. This exercise is not used as a separate

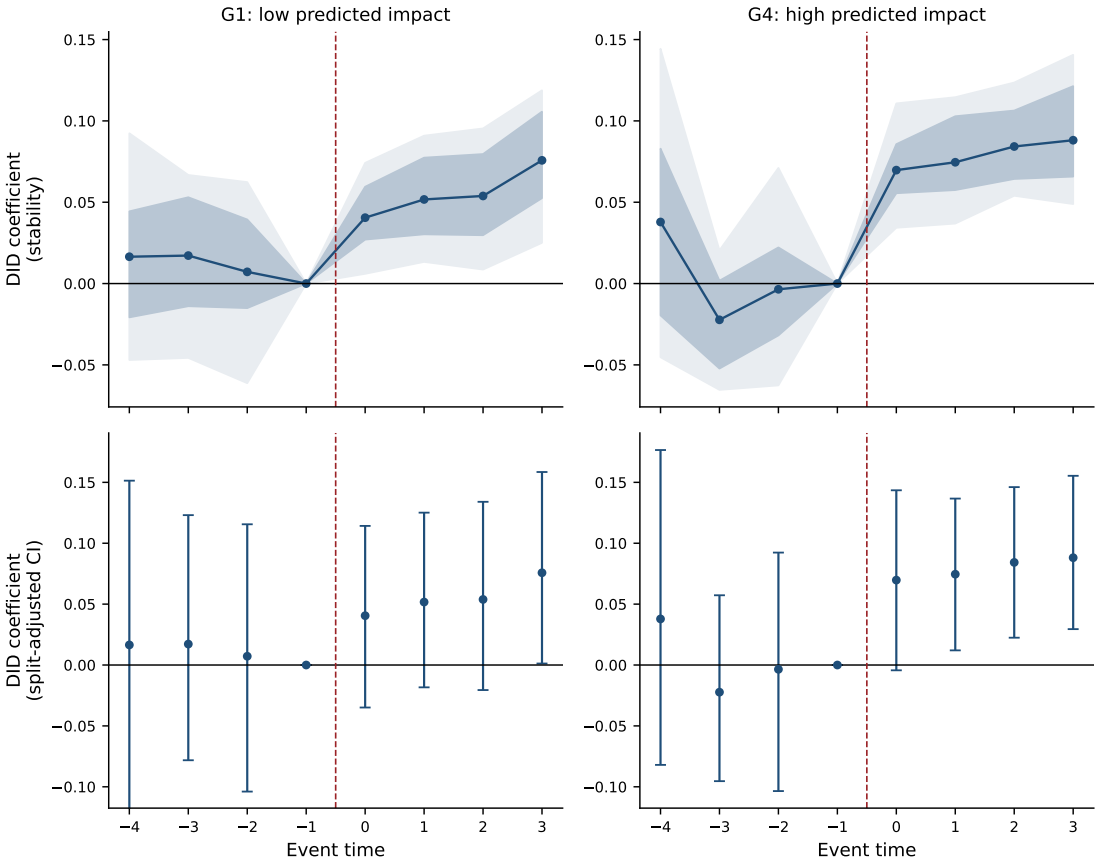
heterogeneity design, because it depends on the subset of households that can be reobserved and transported into the panel sample. Instead, its role is more limited: it assesses whether the profile ranking learned in the repeated cross-section is directionally preserved in a different sample. The point estimates remain ordered in the expected direction, with larger within-household gains among higher-score households. However, because the household-panel validation uses a substantially smaller sample, the group-specific estimates are considerably less precise than the repeated-cross-section GATES estimates. I therefore interpret this exercise as a directional validation of the score ordering rather than as evidence of statistically precise heterogeneous effects.

Figure A.1: Household-panel validation of heterogeneous effects on community freedom



Notes: The figure reports an auxiliary validation in reobserved households. For each split, cantons used for household-panel validation are excluded from score training. The score is trained in the repeated cross-section and applied to reobserved households, where household fixed-effects difference-in-differences estimates are computed by predicted-impact group. In Panel A, blue circles show split-adjusted confidence intervals; red squares show the interquartile range of split-specific point estimates and summarize split stability. The panel validation is intended as a directional check of the score ordering in a smaller reobserved-household sample.

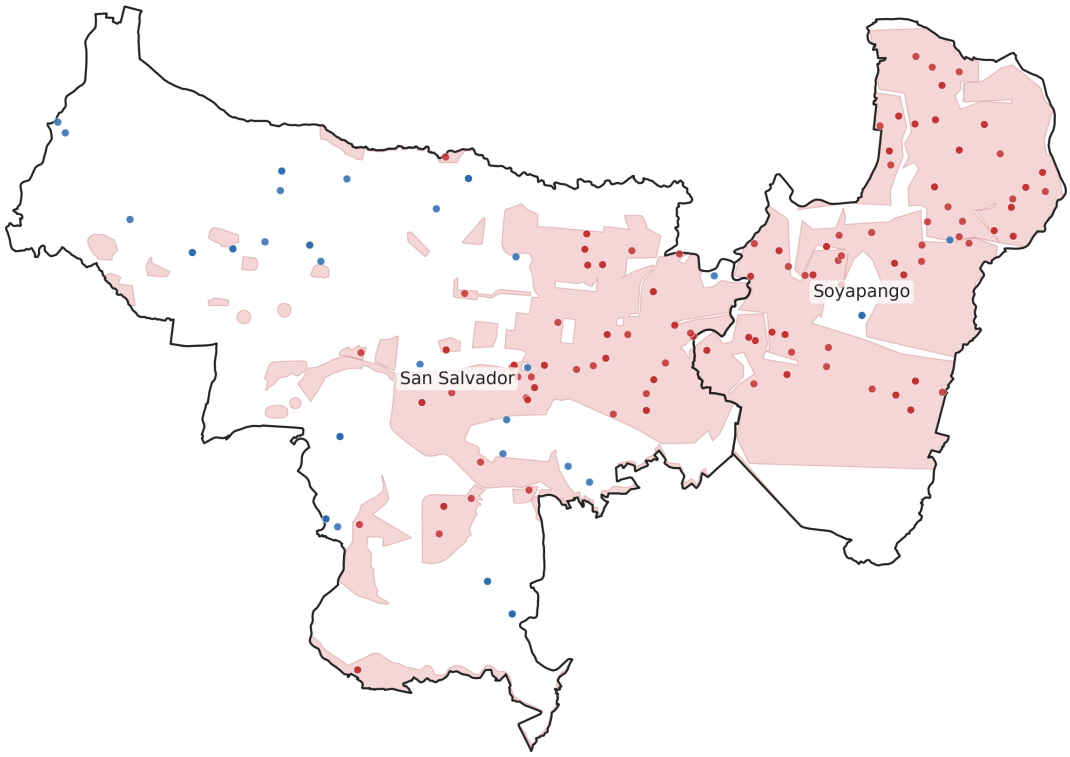
Figure A.2: Event-study diagnostics by predicted-impact group



Notes: The figure reports event-study profiles for the lowest predicted-impact group and the highest predicted-impact group. Event time is measured relative to the beginning of the crackdown in 2022Q2, and the omitted period is the event-year immediately before the crackdown.

A.3 Robustness and Diagnostic Checks

Figure A.3: Spatial validation using gang-control polygons



Notes: The figure compares gang-related homicides in 2016–2020 with an auxiliary shapefile of gang-control polygons for San Salvador and Soyapango. The San Salvador polygons are based on the map from InSight Crime (2019), and the Soyapango polygons are based on a map attributed to El Salvador's National Civil Police and reproduced by Fundación Padre Arrupe (2021). Red shaded areas indicate mapped gang-control polygons. Red points are gang-related homicides that fall inside mapped gang-control polygons; blue points fall outside those polygons. In San Salvador, 81 gang-related homicides fall inside mapped polygons and 50 fall outside. In Soyapango, 103 gang-related homicides fall inside mapped polygons and 4 fall outside. Gang-related homicides outside mapped polygons are, on average, 430 meters from the nearest mapped gang-control polygon.

Table A.2: Respondent Composition: Weighted and Unweighted Estimates

	Survey weighted		Unweighted	
Outcome	$\widehat{\beta}$ (SE)	N	$\widehat{\beta}$ (SE)	N
Female-headed household	-0.002 (0.009)	118,816	0.004 (0.007)	118,816
Household-head age	0.014 (0.319)	118,816	-0.237 (0.238)	118,816
Number of men	0.007 (0.019)	118,816	0.011 (0.014)	118,816
Number of women	0.023 (0.022)	118,816	0.009 (0.017)	118,816
Household size	0.030 (0.033)	118,816	0.020 (0.024)	118,816
Children ages 0–12	-0.006 (0.014)	118,816	-0.001 (0.011)	118,816
Adolescents ages 13–17	0.003 (0.009)	118,816	0.007 (0.007)	118,816
Males ages 12–35	0.013 (0.015)	118,816	0.016 (0.011)	118,816
Any male ages 12–35	0.005 (0.010)	118,816	0.011 (0.007)	118,816
Members ages 65+	0.018* (0.011)	118,816	0.008 (0.009)	118,816
Dependency ratio	-0.000 (0.010)	108,329	-0.001 (0.009)	108,329
Joint Wald test p-value	0.388		0.171	
Common-sample households	108,329			

Notes: The table reports household-level difference-in-differences estimates for observable respondent-composition variables. Survey-weighted estimates use EHPM household weights; unweighted estimates describe the raw respondent pool. The coefficient is the interaction between the post-2022Q2 period and pre-policy gang-related homicide exposure. All regressions absorb canton and year-quarter fixed effects and cluster standard errors by canton. The joint Wald tests use the common complete-case sample. Stars denote $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Table A.3: Dwelling and Reporting Placebos

Outcome	$\widehat{\beta}$ (SE)	Mean	N	Clusters	R^2
Dwelling characteristics and basic services					
Rooms in the dwelling	-0.038 (0.024)	2.765	118,816	746	0.117
Bedrooms in the dwelling	-0.037* (0.021)	1.600	118,816	746	0.109
Electricity as main lighting source	-0.011* (0.006)	0.983	118,816	746	0.068
Piped-water access	0.002 (0.009)	0.817	118,816	746	0.297
Reported receipt of government subsidies					
Reports gas subsidy	-0.012 (0.011)	0.536	118,661	746	0.063
Reports electricity subsidy	0.008 (0.012)	0.403	117,003	746	0.066
Reports basic universal pension	0.005* (0.003)	0.007	118,412	746	0.051

Notes: The table reports survey-weighted difference-in-differences estimates for dwelling characteristics, basic services, and reported receipt of government subsidies. The coefficient is the interaction between the post-2022Q2 period and pre-policy gang-related homicide exposure. All regressions absorb canton and year-quarter fixed effects and cluster standard errors by canton. Mean is the pre-treatment mean among exposed cantons. Gas, electricity, and pension outcomes use harmonized 2018–2019 and 2020–2025 questionnaire items. Stars denote $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Table A.4: Unconditional Use of Recreational and Community Spaces

Outcome	$\widehat{\beta}$ (SE)	Mean	%Δ	N	Clusters	R^2
Sports court used	0.035*** (0.012)	0.276	12.7	118,148	746	0.090
Park or green area used	0.038*** (0.011)	0.142	26.6	118,148	746	0.138
Children's play area used	0.023*** (0.007)	0.086	26.6	118,148	746	0.089
Community center used	0.022** (0.011)	0.126	17.1	118,148	746	0.167
Any recreational/community space used	0.050*** (0.014)	0.345	14.6	118,148	746	0.111
Number of recreational/community spaces used	0.117*** (0.033)	0.631	18.6	118,148	746	0.131

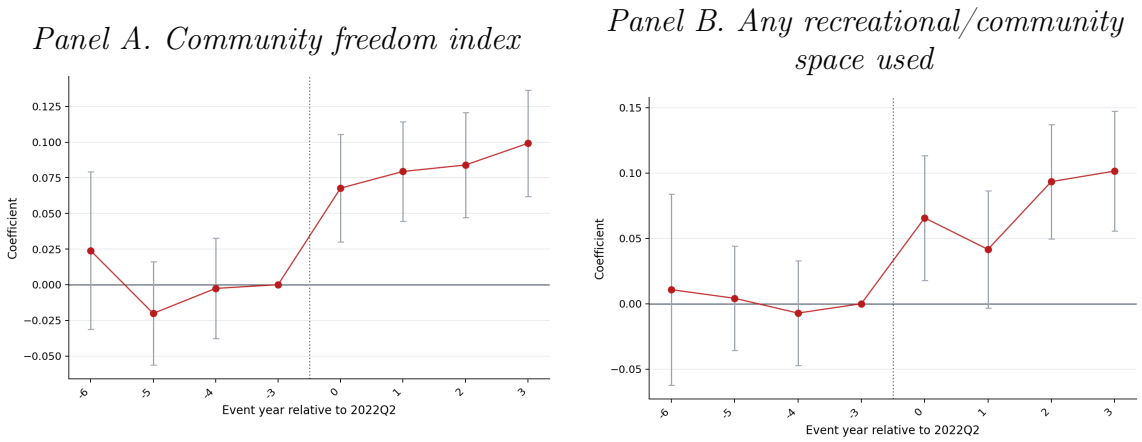
Notes: The table reports survey-weighted difference-in-differences estimates for the outcomes corresponding to Panel B of Table 1, defined without conditioning on reported facility availability. A household is coded as a nonuser when it reports that a facility is unavailable; the outcome remains missing when the facilities module is not observed. The inherently conditional use rate among available spaces is not reported. Regressions absorb canton and year-quarter fixed effects and cluster standard errors by canton. Mean is the pre-treatment mean among exposed cantons. Stars denote $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Table A.5: Split-Year IV Sensitivity to Gang-Exposure Measurement Error

Outcome	Baseline DID (SE)	Median split IV	Median SE	5th–95th pct.	Min. first-stage F
Community freedom index	0.077 (0.013)	0.154	0.031	[0.129, 0.183]	41.4
Any recreational/community space used	0.055 (0.013)	0.117	0.026	[0.094, 0.143]	41.0

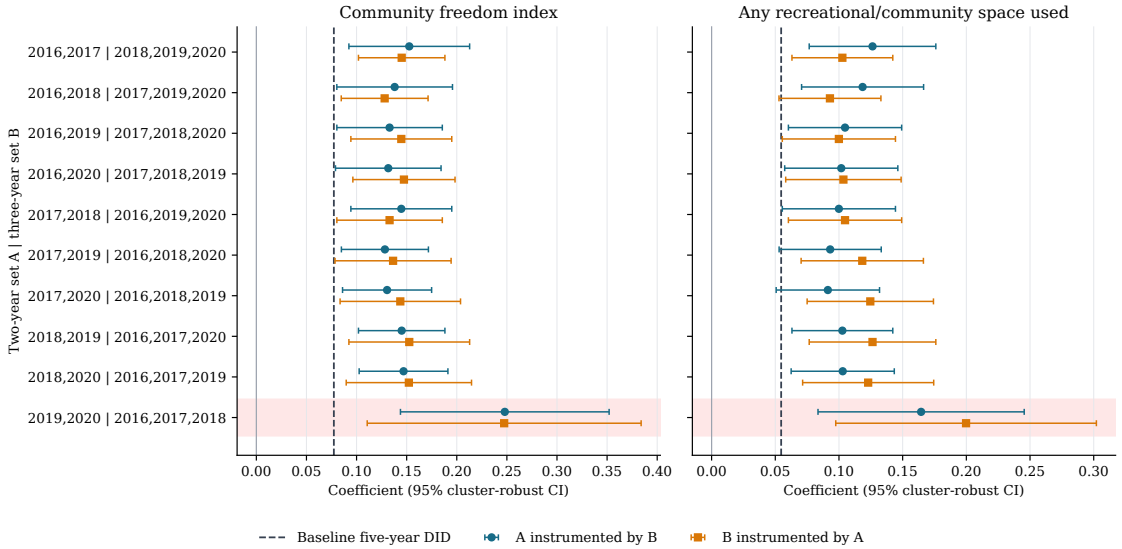
Notes: Each of 500 seeded draws assigns two of the five measurement years to proxy A within each canton and the remaining three years to proxy B. Each proxy is then instrumented by the other, yielding 1,000 estimates per outcome. The 5th–95th percentiles summarize sensitivity to the random assignment and are not confidence intervals. Every split-IV estimate is positive, exceeds the corresponding baseline DID estimate, and has a 95 percent confidence interval that excludes zero. The baseline uses the full 2016–2020 exposure indicator. All specifications use EHPM household weights, absorb canton and year-quarter fixed effects, and cluster standard errors by canton.

Figure A.4: Event-study estimates using a pre-COVID reference period



Notes: The figure reports event-study estimates for the two main public-life outcomes using a pre-COVID benchmark for the pre-crackdown reference period. The specification excludes survey years 2020 and 2021 and 2022Q1, and adds 2017 and 2018 to the pre-period. Event time is measured in event years relative to 2022Q2. The omitted period is the last available pre-crackdown event year constructed from pre-COVID survey years.

Figure A.5: Split-Year IV Estimates Across Common-Calendar Partitions



Notes: Each row partitions the 2016–2020 homicide records into the two-year set A shown to the left of the vertical bar and the complementary three-year set B. Circles instrument the exposure indicator from set A with the indicator from set B; squares reverse the instrument direction. Dashed vertical lines report the baseline DID estimates using the full five-year exposure indicator. The shaded row isolates the relatively sparse 2019–2020 two-year proxy. All estimates use EHPM household weights, absorb canton and year-quarter fixed effects, and cluster standard errors by canton.

Table A.6: Dose-Response Estimates by Gang-Related Homicide Intensity

Intensity group	β	SE	Mean	%Δ	N	Clusters	R^2
Outcome A. Community freedom index							
Low intensity × Post	0.079***	0.017	0.732	10.7	115,509	746	0.164
Middle intensity × Post	0.094***	0.020	0.719	13.1	115,509	746	0.164
High intensity × Post	0.053***	0.016	0.755	7.0	115,509	746	0.164
Outcome B. Any recreational/community space used							
Low intensity × Post	0.064***	0.014	0.507	12.6	78,279	746	0.068
Middle intensity × Post	0.053**	0.022	0.529	10.1	78,279	746	0.068
High intensity × Post	0.023	0.019	0.590	3.9	78,279	746	0.068

Notes: The table reports repeated-cross-section DID estimates from one regression for each outcome listed in the table. Outcome A reports estimates for the community-freedom index; Outcome B reports estimates for use of any recreational or community space. The binary high-exposure indicator is replaced with three post-period interactions for low-, middle-, and high-intensity exposure groups. Intensity is defined by terciles of the 2016–2020 gang-related homicide rate per 100,000 residents among cantons with positive exposure in the main DID canton support. Cantons with zero gang-related homicides in 2016–2020 are the omitted comparison group. Mean is the pre-treatment mean for the corresponding intensity group. All regressions absorb canton and year-quarter fixed effects, use household weights, and cluster standard errors by canton.

Table A.7: Robustness to Urban-Specific Post Trends

	Community freedom index				Any recreational/community space used			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
High exposure × Post	0.077***	0.060***	0.062***	–	0.055***	0.025**	0.025**	–
	(0.013)	(0.012)	(0.012)	–	(0.013)	(0.013)	(0.012)	–
High exposure × Post × Urban	–	–	–	0.072***	–	–	–	0.052***
	–	–	–	(0.022)	–	–	–	(0.017)
High exposure × Post × Rural	–	–	–	0.052***	–	–	–	0.004
	–	–	–	(0.014)	–	–	–	(0.018)
Urban × Post	–	0.039***	–	0.025	–	0.066***	–	0.032*
	–	(0.013)	–	(0.021)	–	(0.013)	–	(0.019)
Urban × year-quarter controls	No	No	Yes	No	No	No	Yes	No
Canton FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	115,509	115,509	115,509	115,509	78,279	78,279	78,279	78,279
Cantons	746	746	746	746	746	746	746	746
R^2	0.164	0.165	0.166	0.165	0.068	0.069	0.070	0.069

Notes: The table reports repeated-cross-section DID robustness checks for the two main public-life outcomes. The coefficient is the interaction between the post-2022Q2 period and an indicator for cantons with at least one gang-related homicide in 2016–2020. Column 1 reports the baseline canton fixed-effects specification. Column 2 adds Urban × Post. Column 3 adds interactions between Urban and each survey year-quarter, omitting 2022Q1. Column 4 allows the high-exposure post effect to differ between urban and rural cantons while controlling for Urban × Post. Urban is an indicator for cantons with weighted pre-treatment urban share at least 0.5. All specifications absorb canton and year-quarter fixed effects, use household weights, and cluster standard errors by canton. Standard errors are in parentheses.

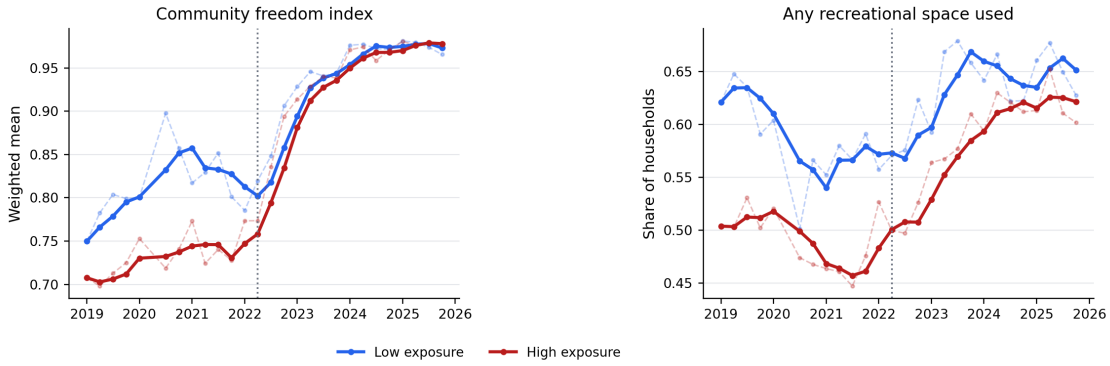
Table A.8: Robustness to Densely Populated Places

	Baseline		No SS/Soyapango		AMSS + SA + SM		Rest of country	
	Freedom	Any Rec. Used	Freedom	Any Rec. Used	Freedom	Any Rec. Used	Freedom	Any Rec. Used
High exposure × Post	0.077***	0.055***	0.068***	0.049***	0.066**	0.140***	0.043***	0.024*
	(0.013)	(0.013)	(0.013)	(0.013)	(0.033)	(0.052)	(0.012)	(0.013)
Mean	0.733	0.526	0.744	0.533	0.673	0.466	0.774	0.576
Change %	10.5	10.4	9.2	9.1	9.8	30.2	5.5	4.2
Clusters	746	746	744	744	89	89	657	657
Obs.	115,509	78,279	113,388	76,582	24,188	16,718	91,321	61,561
R^2	0.164	0.068	0.161	0.070	0.195	0.041	0.145	0.073

Notes: The table reports repeated-cross-section DID robustness checks for the two main public-life outcomes. The baseline columns reproduce the main canton fixed-effects specification. The second pair excludes the two AMSS municipalities that are single-canton in the geographic universe, San Salvador and Soyapango. The third pair restricts the sample to the densely populated-place scope: AMSS-touched municipalities from the gang-polygon validation plus the municipalities of Santa Ana and San Miguel. The final pair reports the complement of that scope. All regressions absorb canton and year-quarter fixed effects, use household weights, and cluster standard errors by canton. Mean is the pre-treatment mean among exposed cantons. Standard errors are in parentheses.

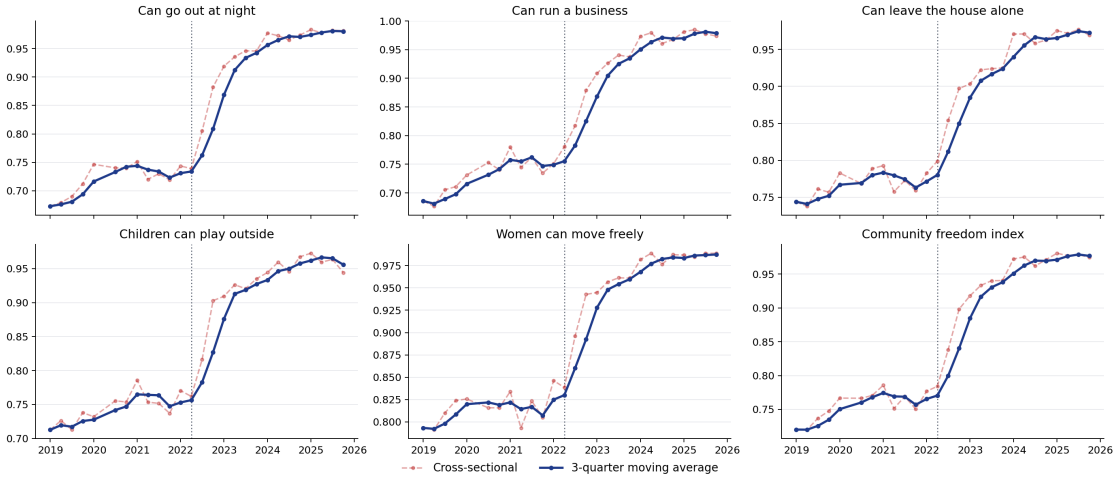
A.4 Additional Descriptive Trends in Local Public Life

Figure A.6: Descriptive trends in key public-life outcomes by exposure group



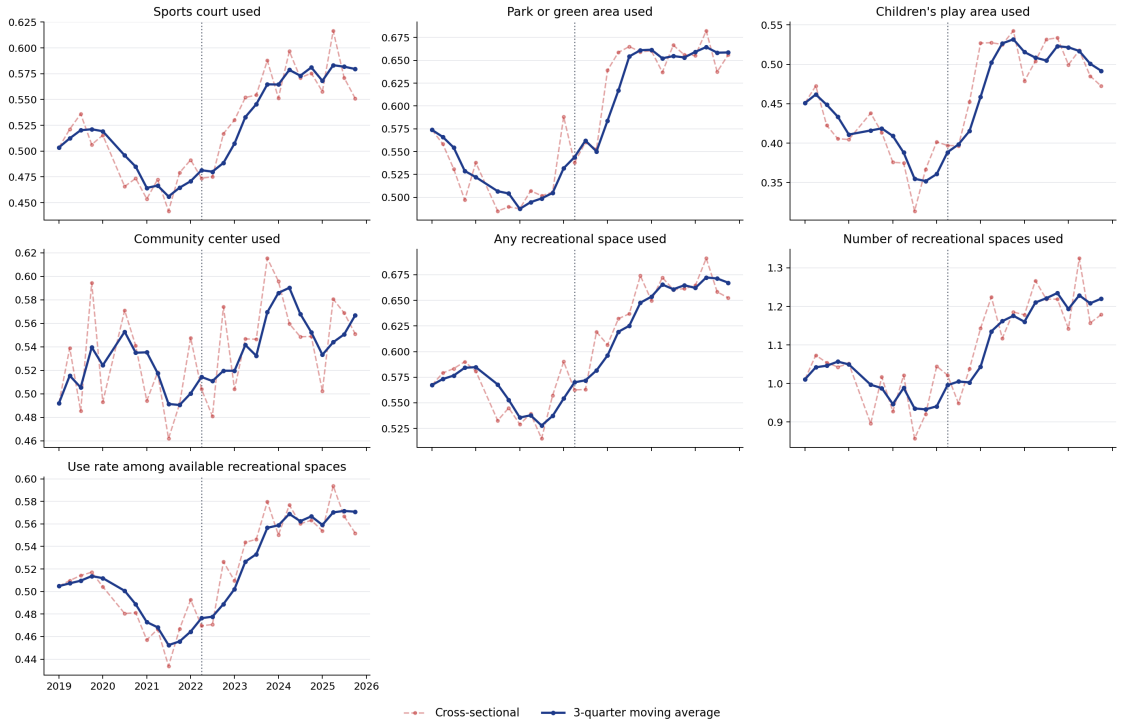
Notes: The figure reports weighted quarterly means for the two main public-life outcomes separately for cantons with and without at least one gang-related homicide in 2016–2020. Solid lines report three-quarter moving averages; faint dashed lines report raw quarterly weighted means. The vertical dotted line marks 2022Q2.

Figure A.7: Descriptive trends in community freedom outcomes



Notes: The figure reports weighted quarterly means for the community-freedom index and the survey items that compose it.

Figure A.8: Descriptive trends in recreational and community space use



Notes: The figure reports weighted quarterly means for use of recreational and community spaces, including each space type and aggregate measures of use.

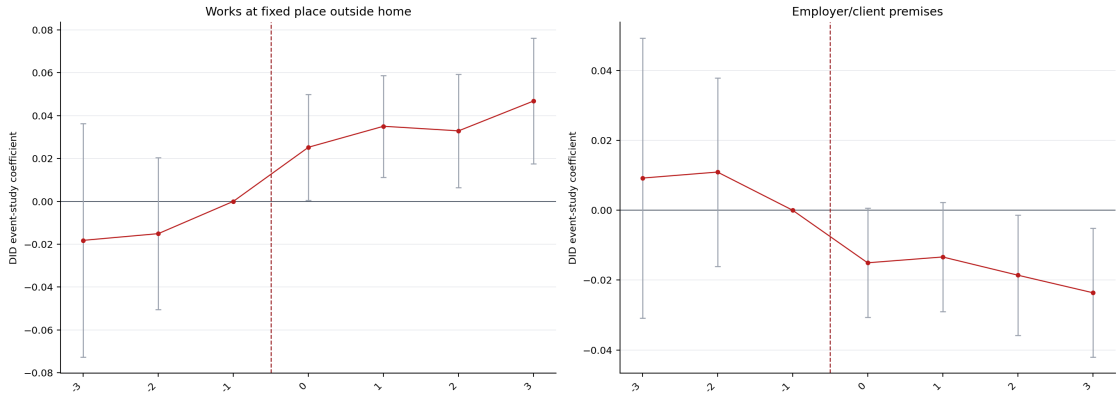
A.5 Exploratory Labor-Market Outcomes

Table A.9: Difference-in-differences estimates for selected labor and productive-mobility outcomes

Outcome	β	SE	Mean	%Δ	N	Clusters	R^2
Panel A. Labor-market participation, earnings, and search							
Employed	-0.004	0.007	0.581	-0.665	256,585	746	0.013
In labor force	-0.002	0.007	0.607	-0.333	256,585	746	0.014
Unemployed	0.002	0.002	0.026	7.154	256,585	746	0.008
Inactive	0.002	0.007	0.393	0.514	256,585	746	0.014
Labor income, employed	12.13	8.023	280.4	4.326	152,948	746	0.084
Searched job or business	0.003	0.005	0.063	4.983	96,905	746	0.023
Panel B. Workplace location							
Works at home	-0.009	0.006	0.119	-7.609	132,297	746	0.020
Works at fixed place outside home	0.042***	0.010	0.693	6.017	132,297	746	0.028
Employer/client premises	-0.022***	0.008	0.069	-32.09	132,297	746	0.025
Mobile or public-facing workplace	-0.010	0.008	0.119	-8.692	132,297	746	0.021

Notes: The table reports difference-in-differences estimates for selected labor-market outcomes among adults ages 15–64. The coefficient is the interaction between the post-2022Q2 period and pre-policy gang-related homicide exposure. Mean is the pre-treatment mean among exposed cantons. Regressions absorb canton and year-quarter fixed effects, use person weights, and cluster standard errors by canton. Labor income is measured conditional on employment. These outcomes are exploratory and are not part of the main mobility result. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Figure A.9: Event-study estimates for selected labor and productive-mobility outcomes



Notes: Panels report event-study coefficients for selected exploratory labor and productive-mobility outcomes among adults ages 15–64. Treatment is canton-level exposure to at least one gang-related homicide in 2016–2020. Event time is measured relative to 2022Q2, with event time −1 omitted. Event time −4 is not shown because the relevant question was not asked or was not asked in an equivalent form. Regressions absorb canton and year-quarter fixed effects, use person weights, and cluster standard errors by canton. Whiskers show 95 percent confidence intervals.