Kingpin Down: Power Vacuums, Market Structure, and the Violent Consequences of High-Profile Arrests *

Heather Bone†
University of Toronto

Monday 2nd June, 2025

Job Market Paper

Click here for the latest version of this paper.

Abstract

This paper studies how targeting the top leaders of criminal enterprises impacts market structure and homicide rates in Mexico. While executive capture may diminish organizational capacity, it may also induce territorial disputes as the existing balance of power between organizations is disrupted. To empirically assess these competing theories, I build a new longitudinal dataset on the presence of criminal enterprises in Mexican municipalities, drawing from over 21 million newspaper articles. I train natural language models to extract (1) whether the article pertains to organized crime, (2) the names of criminal enterprises and locations mentioned in the text, and (3) which locations (if any) each criminal enterprise is operating in. I find that capture or assassination induces a 54% local increase in the number of criminal enterprises operating in these locations, reflecting the entry of competing groups to the geography. This market restructuring explains why a leader's capture or death induces a 32% increase in local homicides, suggesting costly repercussions associated with the policy.

* I would like to thank Arthur Blouin, David Price, and Kory Kroft for their guidance throughout this project. Additional thanks to Taryn Eames, Kate Rybczynski, Marie Harrigan, Robert Martens, Mike Mueller-Smith, Jenna Shelton, Barb Lundebjerg, and Chris Blattman for their helpful comments. I am grateful to participants in the University of Toronto Empirical Microeconomics Brown Bag seminar for their valuable feedback.

† Email: heather.bone@mail.utoronto.ca
Website: https://www.heatherbone.com
Address: Department of Economics, University of Toronto, 150 St. George St, Toronto, ON, Canada

1 Introduction

In much of Latin America, violence related to drug trafficking and organized crime has risen considerably over the past two decades, prompting substantial debate over the best policies to improve public safety (UNODC (2013), UNODC (2019)). In this paper, I study one popular enforcement strategy employed by governments across the world: high-value targeting (henceforth HVT), which is also referred to as the kingpin strategy. Under this strategy, law enforcement targets the top leaders, or kingpins, of criminal enterprises (henceforth CEs), resulting in their capture or death. 14 of the 23 countries classified by the United States government as major drug-producing or transit countries have implemented HVT against kingpins.[1] This study seeks to address two key questions: Does high-value targeting (HVT) affect levels of violence, and if so, through what mechanisms?

HVT has been controversial because the policy has ambiguous effects on violence, depending on how it affects market structure. Supporters of the policy argue that removing a CE's kingpin may weaken the group's organizational capacity and, as a result, cause the CE to collapse, leading to a reduction in violent crime (Molzhan et al., 2014). In contrast, opponents argue that capturing or killing high-ranking leaders may instead exacerbate violent crime by creating a power vacuum if no clear successor is in place. This could induce infighting as members compete for control, potentially causing larger CEs (called cartels) to splinter into smaller groups.[2] Critics also argue that targeting a high-ranking leader may weaken the CE, making it vulnerable to attacks by rivals seeking to usurp its territory (Ríos (2013), Calderón et al. (2015), Osorio (2015), Phillips (2015)).

To design policies that effectively reduce criminal violence, it is not enough to ask if HVT affects violence; we must also understand the underlying mechanisms. If HVT effectively reduces violence by eliminating the targeted group, expanding its use could prove beneficial. If HVT increases violence by causing CEs to splinter into new groups, policies should focus on resolving disputes within CEs to prevent factional conflicts.[3] Conversely, if competition among rival groups is the main driver of violence, demand-side interventions aimed at

[1] These countries are published annually by the White House in the Memorandum on Presidential Determination on Major Drug Transit or Major Illicit Drug Producing Countries. I use data from the 2024 list. Table A1 details, for each country on the list, whether it has used HVT, along with supporting evidence.

[2] While supporters of the policy acknowledge this as a possibility, they claim these smaller organizations are easier for law enforcement to eliminate. Thus, they claim that even if splintering occurs, the policy results in a long-run reduction in violent crime (Molzhan et al., 2014).

[3] For instance, programs that employ former gang members to mediate conflicts have been effective in reducing factional violence, as shown by Avram et al. (2024).

lowering market profitability may be more effective in reducing violence.[4]

I study the effects of HVT within the context of Mexico. Mexico is an ideal setting due to its central role in global drug trafficking, its high rates of violent crime, and its aggressive implementation of HVT. Because it shares a nearly 2000 mile border with the United States, Mexico plays a central role in moving drugs north to one of the world's largest illicit drug markets (UNODC, 2010).[5] Although drug trafficking has a long history in Mexico, as Figure 1 shows, homicide rates only began to surge in 2008, increasing 162% between 2008 and 2019. As a result, Mexico is now one of the world's most violent countries: Eight of the ten global cities with the highest rates of homicides per capita are in Mexico (Seguridad Justicia y Paz, 2022).[6] This surge followed the election of President Felipe Calderón in 2006, who employed HVT at an unprecedented frequency. Calderón's administration captured or killed nine times as many kingpins as his predecessor Vicente Fox, as shown in Figure 2. Although Calderón cited escalating regional violence as justification for employing HVT, national homicide rates had been declining before his administration's shift to this strategy, leading to significant debate about whether HVT inadvertently escalates violence rather than reducing it.[7]

To contribute to the debate, I estimate how market structure and violence are changing in municipalities in which kingpin capture or deaths take place, relative to a counterfactual in which they did not occur. Addressing this empirical challenge is complicated, as it requires panel data that comprehensively identifies both the periods and specific regions in which CEs operate. Researchers commonly use news data to track this information, as it allows them to track more CEs at a more disaggregated geographic level and at a higher temporal resolution (Coscia and Rios (2012), Sobrino (2019), Osorio and Beltran (2020)). Despite this, existing data on market structure is incomplete due to the absence of a definitive list of CEs in Mexico. This is due to the fact that, in recent years, the rapid proliferation of CEs, coupled with their increasingly smaller sizes, has made these organizations more difficult to

[4] Gavrilova et al. (2019) report that legalization of cannabis in U.S. border states has led to declines in violent crime, consistent with reduced black-market profitability leading to less competition-driven violence among criminal groups.

[5] Estimates of the amount of cash sent from the United States to Mexican CEs range from $19 billion USD to $29 billion annually (DHS, 2010)

[6] Cities in active war zones and cities with populations less than 300,000 were excluded from the study.

[7] The high resource demands of HVT are another reason for its widespread criticism. Although the exact costs of HVT are not reported, it formed a central part of the national security strategy during the Calderón and Peña Nieto administrations, placing heavy demands on the national budget. In 2015, Mexico's spending on the military, domestic security, and the justice system exceeded $12 billion; this amount represented 4.6% of Mexico's national budget for that year (Schippa, 2016).

monitor.[8]

To address this challenge, I develop a new dataset that tracks the presence of criminal groups in Mexico between 2005 and 2019 using tools from natural language processing. Specifically, I train a system of semi-supervised deep learning models to process over 21 million Spanish-language newspaper articles from Latin America.[9] These enable me to identify three key features in the raw data: (1) whether the news articles pertain to CEs, (2) which exact CE is referenced, and (3) where geographically they appear to be operating. My approach, which does not rely on pre-defined lists of CEs, enables me to track the presence of small criminal organizations that increasingly characterize Mexico's criminal landscape and are potentially important in terms of explaining violent outcomes. As a result, I am able to track the presence of 408 unique criminal enterprises at the municipality-year level. The largest number tracked by other researchers is 171 (Osorio and Beltran, 2020). I then pair the CE panel with data on kingpin captures and deaths derived from the list of sanctions made under the United States Foreign Narcotics Kingpin Designation Act. The list publicly identifies significant foreign narcotics traffickers, including Mexican kingpins. To study the effects of HVT on violence, I use quarterly data on homicides per 100,000 population, which I derive using death certificate data from the Mexican health ministry and population data from the Mexican census.

My dataset allows me to uncover stylized facts about the structure of Mexico's criminal landscape. These facts help contextualize the environment in which HVT was implemented. In particular, I reveal that the market structure remains concentrated within key municipalities over time. Within these municipalities, competition among CEs is growing. This is driven primarily by the presence of smaller, less prominent criminal organizations that were previously difficult to track. By developing these stylized facts, I can explore patterns of competition and fragmentation at a level of detail that has not been captured in earlier work.

To test how HVT affects market structure and homicide rates in municipalities where kingpin captures and deaths occur, I use the synthetic difference-in-differences (SDID) estimator, which utilizes spatial and temporal variation in the policy's implementation to iden-

[8] For example, an investigation by Vela (2020) found that in Mexico City alone, the number of CEs grew from 25 to 40 within a single year. Reflecting these challenges, Stratfor, a strategic intelligence publisher known for its annual maps of Mexican CE territories, has limited its analysis to select regions and now focuses only on the most prominent CEs (Stewart, 2019).

[9] The models are semi-supervised because they use labelled training data in addition to using unsupervised models to embed context in vector representations of text. They are trained and evaluated on data from the ProQuest Latin American Newsstream.

tify its effects.[10] SDID draws inspiration from both the standard difference-in-differences and synthetic control estimators, thus incorporating advantages from both models. Specifically, for each treatment date, SDID constructs a synthetic control unit by weighting never-treated municipalities to match the pre-treatment trends of municipalities where kingpins were captured or killed. The SDID procedure inherently accounts for unit fixed effects to address time-invariant differences across municipalities and time fixed effects to capture shocks common to all units in each period. In this paper's context, where the characteristics of municipalities experiencing kingpin captures differ from those of most never-treated municipalities, the synthetic control method improves validity by selecting a subset of units that align more closely with the treated municipalities' historical trends.

Using this approach, I find that the average kingpin capture or death leads to a 54% increase in the number of locally operating criminal enterprises. While pre-capture trends are comparable across municipalities, the sudden surge in the number of groups post-treatment suggests a market disruption. To further investigate this mechanism, I distinguish between new CEs entering municipalities and splinter groups that have broken off from existing organizations. Specifically, I estimate the policy's effect on the number of non-new splinter groups that existed prior to treatment versus new splinters. My results show that the increase in criminal presence is primarily driven by new entries, consistent with the idea that rival groups are exploiting the weakened position of incumbents to expand their market share.[11]

These results align with the economic model proposed by Osorio (2015), which posits that HVT weakens the organizational capacity of the incumbent criminal enterprise. By reducing the incumbent's ability to defend its territory, HVT lowers barriers to entry, creating a more contestable market where rival groups are incentivized to expand their operations. Though the number of groups eventually declines from its peak, it remains persistently higher many years after the policy is implemented, and there is no evidence that existing groups leave the market. This suggests a permanent decrease in market concentration as new groups enter to compete with incumbents.

The results also imply that violence should increase in response to elite capture. Consistent with this, I find that kingpin capture or death causes the local homicide rate to increase by 32% in the four years following treatment. That competition between CEs is often violent reflects the fact that in illicit markets, there are no formal dispute mechanisms between competitors (Levitt and Venkatesh, 2000). Furthermore, I find evidence consistent

[10] Specifically, since kingpin captures and deaths occur in different municipalities at different times, I follow Clarke et al. (2023), who extends the estimator developed by Arkhangelsky et al. (2021) to settings with staggered treatment.

[11] While the results for splintering are economically significant, they are imprecise and not statistically significant. While I cannot rule out splintering, the estimates are much smaller than those for new entry.

with HVT-related violence spilling over to the general population, with homicides increasing in both demographic groups more likely to be involved in criminal enterprises and those not typically associated with such activity. These findings highlight the significant welfare costs of targeting the leaders of criminal enterprises, as the resulting violence affects not only criminal actors but also the broader population, undermining public safety and stability in affected municipalities. This is consistent with prior research emphasizing the civilian costs of criminal violence in Mexico, which include shifts in risk-taking behavior, disruptions to employment, and increases in migration (Utar (2018), Orozco-Aleman and Gonzalez-Lozano (2018), Brown et al. (2019)).

Measurement error in the CE panel, whether due to variation in reporting over time or threats against journalists could, in principle, bias my results. If reporting bias is both correlated with treatment and evolves differently over time between treated municipalities and their associated synthetic controls, it distorts the estimated treatment effect, as systematic differences over time confound the results. To address this, I compare my data with the Mapping Criminal Organizations (MCO) Project. Unlike newspaper-based data, which can be affected by local intimidation or shifts in media presence, the MCO data leverages government and expert sources that are less susceptible to these biases. Instead, these sources draw on intelligence networks and official investigations.[12] By comparing my newspaper data with the MCO data, I can gauge how reliably each CE's presence is documented. The results show that the most prominent CEs, which might be more likely to suppress coverage, are consistently captured in both datasets. Additionally, the reporting trends are largely parallel across high- and low-population states, strengthening the credibility of my results by reducing concerns that observed effects are driven by inconsistent reporting across regions and time. As further evidence, I also use the volume of newspaper coverage that does not relate to organized crime as a dependent variable in my model and find no evidence of differential increases in this coverage that aligns with treatment timing.

To address the risk of positive selection, I rely on detailed narratives for each event to exclude kingpin captures likely influenced by endogenous factors, using an approach inspired by Romer and Romer (2010).[13] I use reporting on captures to exclude cases that could plausibly generate anticipatory responses or were affected by targeted operations in response

[12] While some sources included in the MCO project utilize news data, I only utilize observations that require agreement between all sources to ensure that non-news data is not being used for the comparison.

[13] Romer and Romer (2010) estimate the macroeconomic impact of tax changes, addressing the endogeneity concern that tax policy often responds to economic conditions, which could bias estimates. To isolate the causal effects of tax policy, they construct a series of exogenous tax changes by using a narrative approach, drawing on historical documents such as presidential speeches, economic reports, and congressional records to determine the motivation behind each tax change. By identifying changes driven by long-term goals rather than responses to current economic conditions, they aim to reduce omitted variable bias and obtain more reliable estimates of tax effects on output.

to civilian massacres.[14] The main results are robust to the inclusion of all kingpins.

This research contributes to three main strands of literature. First, this paper adds to the literature studying the effects of increased enforcement efforts against criminal enterprises. Most of this literature documents unintended consequences associated with drug interdiction (Dell (2015), Castillo et al. (2020), Sviatschi (2022)).[15] These papers provide causal evidence of what is commonly referred to as the balloon effect, in which increased levels of drug interdiction causes drug trafficking to flow to areas with less resistance. As a result, child labour and violence increase in these new areas. One notable contribution is Lindo and Padilla-Romo (2018), which studies the causal effects of HVT on homicide rates in Mexico, focusing on the first five captures of high-ranking leaders under President Calderón. Like them, I find that the policy increases homicide rates. However, the primary contribution of this study lies in revealing the mechanisms behind this violence—a crucial insight for policymakers, as different effects necessitate distinct policy responses.[16]

My paper also relates to the literature on the determinants of CE entry and exit into particular markets and its associated effects on violent crime (Sobrino (2019), Hurtado (2023), Alcocer (2023), Battiston et al. (2022), Cruz and Torrens (2023), De Haro Lopez (2023)). Most of these papers demonstrate that changes in the relative profitability of controlling a particular territory can induce CEs to violently compete for territorial control. While these papers utilize positive demand shocks to study these effects, I provide evidence that policies that change the existing balance of power between organizations can also induce turf war by lowering the costs of establishing territorial control.

Finally, this paper relates to the literature that uses natural language models to detect the areas of operation of Mexican criminal enterprises (Sobrino (2019), Osorio and Beltran (2020)). My paper contributes to this literature by employing a named entity recognition model in the data processing pipeline to identify CEs in unstructured newspaper articles. Previous researchers queried the names of CEs and locations in newspaper articles, then used natural language models to predict whether the locations and CEs were linked. In contrast, my model uses contextual clues to identify CEs in unstructured text data, which allows me to track significantly more CEs than previous researchers.

The rest of the paper is organized as follows. In the next section, I discuss the structure of Mexico's criminal landscape and the Mexican government's use of HVT, including the institutional and historical context for my analysis. In Section 3, I present the data and the

[14] Specifically, this includes cases with recent failed capture attempts, false government reports of a kingpin's death, or captures conducted as part of Operation Northern Lynx following civilian massacres.

[15] This includes manual or chemical eradication of drug crops or drug seizures during transport.

[16] As a secondary contribution, this study re-examines the impact on violence by analyzing more than twice the number of events over an extended time period, thereby enhancing the generalizability of the results.

new stylized facts on Mexico's criminal landscape I derive using the new panel dataset. In Section 4, I discuss the empirical strategy and specification. Section 5 presents the results, and Section 6 concludes.

2 Institutional Details

2.1 The Market Structure of Mexico's Criminal Landscape

Mexico's criminal landscape features 3 main classes of groups: cartels, factions, and independent organizations. Cartels are the largest in size and have affiliate organizations, known as cells or factions, that provide services to the larger organization (e.g. protection or money laundering services) (United States Department of Justice, 2010).[17] This structure has become ubiquitous in Mexico, with factions occasionally becoming cartels in their own right. This was the case with one of Mexico's most prominent cartels, Los Zetas, which began as a mercenary wing for the Gulf Cartel and now has factions of its own. Finally, there are independent organized crime groups with no known affiliation, such as small street gangs. By any measure, these organizations play a significant role in Mexico's economy. Collectively, criminal enterprises estimated to employ between 160,000 to 185,000 people in Mexico and earn at least $6 billion in annual profits (Prieto-Curiel et al. (2023), McDonnell (2023)).

Competition between CEs is often violent and is frequently caused by disputes for control of the territory, or *plazas*/turf, that CEs use to conduct their activity. Violence stemming from disputes over territory is commonly called turf war. Ríos (2013) argues that for traffickers, sharing territories increases the costs of corruption, reduces the share of the local market that it can supply, and makes production inputs scarce, thus incentivizing incumbent CEs to invest in military assets to protect their turf and discourage entry. Furthermore, when a CE controls a plaza, they can operate a protection racket by charging their competitors a *piso*, or toll, to traffic drugs through their territory, or by extorting businesses or individuals (Ríos (2013), De Haro Lopez (2023)). For example, the Tijuana Cartel, which controls the most important border crossing into California, charges a tax for shipping drugs through the areas it controls (Jones, 2023). These fees increase the returns from controlling a strategic area, particularly if the CE has a monopoly, as competition diminishes a cartel's ability to charge for protection in this way (Ríos, 2013). Opponents of HVT argue that the policy disrupts the balance of power within targeted organizations, which incentivizes their rivals to try to usurp their territory, as the cost of engaging in this violence falls.

[17] It is worth noting that the use of the word cartel is a misnomer, as Mexican CEs do not currently collude to fix prices or supply and violently compete with each other (Keefe, 2012). The origin of the term is discussed in Section 2.2.

2.2 Historical Context: Pre-Calderón

Mexican CEs have been active for over a century and emerged to traffic alcohol to the United States during alcohol prohibition (Medel and Thoumi, 2014). However, the enforcement strategies employed against them and the market structure of organized crime has evolved over time due to diplomatic pressure from the United States and changes to Mexico's political institutions stemming from the country's democratization in the 1990s. A discussion of this history is necessary to explain why Calderón dramatically expanded Mexico's use of HVT at the beginning of his term. Furthermore, this historical context helps explain why in Mexico, the timing of kingpin captures was plausibly exogenous, making it an ideal setting in which to study the effects of HVT.

From the 1929 to 1988, Mexico was ruled by the Institutional Revolutionary Party (PRI) at all levels of government. During its rule, the relationship between CEs and the state was characterized by corruption. CEs were granted exclusive state-sponsored protection in a territory known as a *plaza* in exchange for sharing their profits with authorities and refusing to commit violence against civilians and state officials. Failure to meet these conditions led to the traffickers' arrest (Astorga (2001), Lupsha (1991)).

At first, these agreements were made with local authorities such as mayors, police chiefs, and large land owners (Lupsha, 1991). This changed in 1947, when the decentralized system became federalized following the establishment of the Federal Security Directorate (DFS), a political police force that took the lead on drug enforcement (Lessing, 2017). Because agreements were now made with a central authority, the need for high-level political connections increased, which favoured politically sophisticated and connected traffickers and promoted market consolidation (Lupsha, 1991).

Enforcement intensified in the late 1960s after the United States worked with Turkish and French law enforcement to disrupt heroin trafficking routes. This created a surge in the demand for Mexican heroin and lead to a rapid expansion of the market, which further favoured large traffickers due to economies of scale in drug production (Lupsha, 1991). In response to increasing drug flows across the southwest border, the Nixon administration launched Operation Intercept in 1969. The operation was a temporary initiative that nearly closed the U.S.-Mexico border by requiring a thorough inspection of all vehicles (Cedillo, 2021). Afterwards, the Mexican government began to collaborate with Washington on an anti-drug campaign that used the military to eradicate drug plantations, interdict drug shipments, and arrest drug traffickers (Medel and Thoumi, 2014). Still, the state-sponsored protection agreements remained in place. To appease the United States Government, Mexican authorities increased enforcement against groups without these agreements, promoting

further consolidation and giving rise to cartels (Lupsha, 1991). The term cartel became common during this period of market consolidation due to the high market share of these CEs. The term was first used in the 1980s, when they became intermediaries for trafficking Colombian cocaine into the U.S. market, despite the absence of collusion with competitors to fix prices (Fraioli, 2020).

The relationship between cartels and the state changed following Mexico's democratization. Following a series of democratic reforms, opposition parties emerged and the PRI began to lose power, first locally throughout the 1990s and then nationally in 2000. As a result, agreements between CEs and governments began to break down and state enforcement against CEs increased (Osorio, 2012).[18] The collapse of these agreements helped fuel an increase in turf war because their dissolution removed the credible threat of enforcement in response to violence (Lessing, 2017). Furthermore, without state protection, cartels built private militias to defend their territories and went to war with each other (Trejo and Ley, 2018). This turf war, however, was localized in nature. Figure 1 shows that at a national level, homicide rates were declining for most of the 1990s and throughout the term of President Vincente Fox (2000-2006).

2.3 High Value Targeting

My data begins in 2005, prior to the end of Fox's term in 2006. Given that overall homicide rates were declining at this time, the public saw economic problems as far more pressing than security-related issues. However, this began to change under President Felipe Calderón, who, shortly after taking office, launched a massive crackdown on cartels, with a focus on removing the kingpins of these organizations. While Calderón used ongoing turf war to justify this offensive, the evidence suggests that the policy was driven primarily by political considerations (Lessing, 2017).

Calderón had won the election by a slim 0.58% margin and his refusal to endorse a recount, coupled with his opponent declaring the result invalid, sparked mass outrage and social unrest. Calderón needed to quickly shift debate away from the election results (Lessing, 2017). For this, he looked to Colombia and, while he was still president-elect, met with Colombian President Àlvaro Uribe, who had an approval rating of over 80%. This meeting inspired Calderón to employ HVT against the cartels, drawing inspiration from Colombia's use of HVT, which was widely viewed as successful (Jiménez, 2015). After Colombia em-

[18] Osorio (2012) argues that this is because agreements between state actors and criminal enterprises are more stable in non-democratic settings due in part to a hierarchical chain of command that coordinates agreements across different levels of government.

ployed the policy, the large cartels splintered into smaller CEs and violence stemming from the drug trade fell (Lessing, 2017).

However, Lessing (2017) argues that Mexico's implementation of the policy was unlike Colombia's in one key way: Mexico's use of HVT was by and large, unconditional on rates of violence.[19] In Colombia, only the most violent leaders were targeted. Mexico, on the other hand, launched what was described by Calderón as a "war without quarter" in which leaders were typically captured or killed as soon as intelligence became available. This mitigates concerns about selection bias when studying the policy's effects on violence, making Mexico an ideal context for this analysis.

Mexico's divergence from Colombia's approach stemmed from two distinctive features of Mexico's institutional context. First, Mexico's history of state-sponsored protection of criminal enterprises under PRI rule led Calderón to fear that targeting only the most violent groups could create the perception that he was selectively protecting specific CEs. In fact, when the director of Mexico's intelligence agency, Guillermo Valdés, questioned the government's policy of "indiscriminate arrests" in front of a public audience with Calderón present, Calderón responded by interrupting and rebuking him (Lessing, 2017).

Second, there is limited coordination amongst law enforcement agencies in Mexico. Each agency has its own hierarchies, protocols, budgets and intelligence sources. As Lessing (2017) argues, a central authority that coordinates actions in response to cartel violence is needed for a conditional response. Instead Calderón distributed a list of priority targets to all law enforcement agencies. They were encouraged to strike against their targets as soon as possible, which they were subsequently rewarded for through media-friendly "perp walks".

Although Mexico's use of HVT was by and large, unconditional, one notable exception exists. Following two brutal civilian massacres carried out by Los Zetas in 2011, Calderón changed his approach, with a former high-ranking intelligence officer reporting that Calderón privately expressed the need to respond to the massacres by targeting the Zetas (Lessing, 2017). Calderón then launched Operation Northern Lynx in the Zetas' strongholds of San Luis Potosí, Coahuila, Nuevo Léon, and Tamaulipas in 2011 (Ramos, 2011). To account for possible selection bias, I exclude captures in these states from the onset of the operation to the end of Calderón's term from my analysis, as detailed in Section 3.3. The event window studied also includes captures and deaths that occurred under Calderón's successor, Peña Nieto, as kingpin captures and deaths continued, and even increased, under his presidency.[20] Because Nieto did not clearly formulate a policy shift and there is no evidence to suggest

[19] Lessing (2017) argues that the reason in which violence fell in Colombia and rose in Mexico is because Colombia's conditional application of the policy incentivized cartels to eschew violence.

[20] See Figure 2.

that Mexico's unconditional approach changed under his leadership, I include captures and deaths from his presidency in my analysis (Lessing, 2017).

3 Data

To evaluate the consequences of a kingpin capture or death, I need three main data sets. First, I need panel data detailing where and when different types of criminal enterprises are operating. Next, I need data on the dates and locations of the captures and deaths of DTO leaders. Finally, I need data on homicides to evaluate the impact of these events on rates of violence. This data is described in detail in the subsequent sections.

3.1 Criminal Enterprise Data

3.1.1 Model Overview

There are 3 comprehensive longitudinal datasets that track where criminal enterprises in Mexico operate at the municipality-year level. All 3 of them use newspaper data. However, all have limitations that make them inappropriate to use in my context. Coscia and Rios (2012) generate queries containing cartel and municipality names and scrape the resulting articles using Google News between 1990 and 2010, aggregating their results to the annual level. They then use probability-based methods to reduce noise in the dataset.[21] Sobrino (2019) updates the data to cover the period from 1990 to 2016 and queries all articles containing a municipality-cartel pair from a set list of 9 cartels. Sobrino improves upon the methodology of Coscia and Rios (2012) by employing a natural language model to predict whether or not a cartel that appears alongside a municipality operates within it. Osorio and Beltran (2020) also employ natural language models to obtain data on the activities of criminal organizations between 2000 and 2018. This data differs from the aforementioned datasets in that it focuses on extracting information on *event data* (defined as a discrete description of someone doing something to someone else in a given time and place based on explicit information mentioned in the text). By focusing on events, they are discarding data that could be used to identify criminal presence, thus, this dataset is not ideal for tracking where CEs operate.

Unfortunately, tracking where criminal enterprises operate has only become more difficult over time as their number has grown and the groups have become smaller. In Mexico City alone, the number of police-identified criminal groups increased from 25 to 40 between 2019

[21] The Google News API is now obsolete, and Google employs technology to prevent Google News from being scraped.

and 2020 (Mexico News Daily, 2020). By comparison, across all of Mexico both Coscia and Rios (2012) and Sobrino (2019) track just 9 groups at the municipality level. While Osorio and Beltran (2020) track 171 groups, this data's focus on violent events is too narrow for my project, necessitating original data collection.

To obtain panel data on where and when criminal organizations operate, I train and apply three machine learning models to extract this information from a large database of Latin American newspaper articles. This section focuses on how I obtain and validate the data, and Section 3.6 derives new stylized facts on the structure of Mexico's criminal landscape using it. Using Spanish-language data from the ProQuest Latin American Newsstream, spanning from April 2004 to November 2022, I access over 21 million full-text articles, which cover local, national, and regional news.[22] This resource allows me to access content restricted by paywalls that would otherwise be inaccessible through web-scraping.[23] While the database includes articles from across Latin America, I focus on Mexican data by mapping text references to geographic coordinates and excluding non-Mexican locations.

I process the data using natural language processing (NLP) tools to extract structured information from unstructured text, like news articles. These models convert text into numerical vectors, capturing the relationships between words in context. For instance, words like "cartel" and locations are represented in ways that reflect their co-occurrence. By using labeled training data, the models learn patterns, enabling them to classify new articles and extract key entities such as criminal enterprises and locations.

My data processing pipeline consists of four stages: text classification, named entity recognition, relation extraction, and geoprocessing/data cleaning. I describe each stage in detail in Appendix B. First, I train a text classifier using hand-labeled data from 5021 articles to predict whether an article is about organized crime. The classifier employs an ensemble approach that combines a convolutional neural network (CNN) and a bag-of-words model.[24] I apply the trained classifier to the entire dataset, then sample 1001 articles predicted to be about organized crime. I label the locations and CEs in these articles and use this data to train a named entity recognition (NER) model that extracts these features from the text using a transition-based parser, which iteratively classifies entities.

Next, I train a relation extraction model to determine whether a criminal organiza-

[22] While I use data from all years to train the model, my figures do not use post-2019 data as this is outside of the observation window.

[23] Many large Mexican newspapers, such as Reforma and El Universal, restrict access to their content through paywalls.

[24] A convolutional neural network (CNN) captures local patterns in the text by processing words in neighboring windows, allowing it to incorporate context by recognizing how a word's meaning changes based on its surrounding words. The bag-of-words model, meanwhile, represents word frequency without regard to order.

tion operates in a specific location, focusing on individual sentences.[25] This model uses a transformer model to embed each word in its context and predicts the relationship using a feed-forward neural network with a logistic activation function. In simpler terms, the transformer captures how each word relates to the words around it, while the neural network uses this information to predict whether the criminal organization is linked to the location mentioned in the text. Finally, I apply a geoparsing algorithm to disambiguate location names and map them to municipalities in Mexico, followed by data cleaning to ensure the distinct identification of criminal enterprises and presence for the same CE between missing years.[26]

I then classify each CE in the panel as either a cartel, a faction, or an independent group. Specifically, in my context, a faction is a CE that is affiliated with or provides services to a cartel. A cartel is any non-faction that at some point in their tenure, had factions of their own. All other CEs are classified as independent groups. As discussed in Section 2.1, factions can change alliances or break from their parent organizations. To account for this, I perform extensive research on each organization in my data and record if it became independent from their main organization or switched alliances, and if so, what year this occurred in. I call this the *splinter date*, which I use to build a time series panel.[27] Because there are notable cases in which the larger organization still exists, but factions are reported as being at war with eachother, I count organizations that are not reported as becoming independent as splinters as well, with the date of their splinter being the day that they first appear in the data (anywhere in Mexico).[28] In Section 5, I use this information to help me distinguish between the two possible mechanisms leading to entry of criminal enterprises: Changes in criminal presence due to splintering and changes in criminal presence due to new entry.

[25] Relation extraction models are more effective at predicting relationships when the entities are within the same sentence, as they struggle to predict relationships when entities are far apart in the text.

[26] I assume that if there are 3 years or less of missing data two observations where a specific criminal enterprises is identified as operating in a specific location, then they operate in that location in the years in between as well.

[27] I use data from newspapers as well as well as reports from subject-matter experts, such as the Mapping Criminal Organizations project. In cases where a change in criminal status was reported, CEs were classified as factions of their parent organization prior to the year, and independent CEs or cartels after, depending on whether or not they have factions of their own at any point in the panel. In some cases, groups classified as factions have factions of their own (which I refer to as subfactions). For example, Los Negros is a faction of Beltran Leyva that originally was allied with the Sinaloa Cartel prior to Beltran Leyva's split with the Sinaloa Cartel. In these cases, the subfaction is considered a faction of their parent organization's cartel while their parent organization is allied to it if they existed when their parent organization splintered (after which they are a faction of their parents).

[28] One prominent example is the feud between Los Chapitos and Cartel Del Mayo, which are both factions of the Sinaloa Cartel

3.1.2 Model Validation

When building machine learning models to generate data on the presence of criminal enterprises, both internal and external validation are crucial for assessing model performance and ensuring the reliability of the results. Internal validation assesses the model's performance on a test set, which consists of 20% of the data that is held out during training. This helps evaluate how well the model can predict criminal presence based on patterns it has learned, without being exposed to this specific portion of the data during training. External validation, meanwhile, compares the model's outputs to independent datasets, ensuring that the results align with real-world observations and can be trusted beyond the training data.

Validation Metrics

When validating natural language processing (NLP) models, particularly for tasks that involve extracting information specific to organized crime—such as determining whether an article is about organized crime, identifying criminal organizations and locations, and determining relationships between them—metrics like *precision*, *recall*, and the *F1 score* are more informative than overall accuracy. These metrics focus on how well the model captures the particular information relevant to the analysis, ensuring that the key entities and relationships are extracted accurately.

Accuracy refers to the proportion of correct predictions out of the total predictions and is calculated as:

$$\text{Accuracy} = \frac{\text{True Positives} + \text{True Negatives}}{\text{Total Observations}}$$

While accuracy is a common metric, in tasks like named entity recognition (NER) or text classification, accuracy may be misleading if the dataset is imbalanced. For instance, if the majority of articles in the dataset are not about organized crime, a model could achieve high accuracy simply by predicting that no articles are relevant. However, this would not address the specific sub-tasks we care about, such as correctly identifying whether the article discusses organized crime or identifying the criminal organizations and locations within it.

Precision is particularly useful for evaluating sub-questions like *Is this article about organized crime?*, *Is this location related to organized crime?*, or *Is this relationship between a criminal organization and a location valid?* Precision measures the proportion of correctly predicted positive instances (e.g., relevant articles or relationships) out of all instances that the model predicted as positive:

$$Precision = \frac{\text{True Positives}}{\text{True Positives} + \text{False Positives}}$$

In the context of the overall model, precision answers the question: *Of all the criminal organizations or relationships predicted by the model, how many were accurately identified?* High precision ensures that the model avoids false positives. Additionally, in sub-tasks like determining whether the article is about organized crime, high precision would mean that when the model flags an article as relevant, it is very likely to be accurate.

Recall, by contrast, measures how well the model captures all relevant positive instances, answering questions like *How many of the articles about organized crime were actually detected?* or *How many of the correct relationships between a criminal organization and a location were identified?* Recall is defined as:

$$\text{Recall} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Negatives}}$$

High recall ensures that the model captures as many relevant instances as possible, reducing the likelihood of false negatives. This is particularly important when using newspaper data, where we may be concerned about systematic underreporting. High recall helps mitigate these concerns by ensuring that the model captures a more comprehensive view of the presence of criminal enterprises.

The *F1 score* is the harmonic mean of precision and recall, combining both into a single metric:

$$F1 = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}$$

The F1 score is particularly useful in imbalanced datasets, where the distribution of positive and negative cases (e.g., articles about organized crime vs. those not about organized crime) is skewed. In such cases, a model could achieve high precision by correctly identifying a few positive instances while missing many others (low recall), or it could identify most positive instances but introduce too many false positives (low precision). The F1 score provides a balanced evaluation of the model's performance by combining both precision and recall into a single metric. In many applications, this balance is what researchers are most concerned about, as it ensures that the model is not only identifying relevant entities or relationships but doing so with minimal errors. As a result, the F1 score is often the key metric reported in papers because it provides a more holistic view of the model's ability to capture relevant information without being misled by imbalances in the data.

Internal Validation

The internal validation metrics are reported in Table C1.[29] The text classification, named entity recognition (for locations and criminal enterprises), and relation extraction models have F1 scores of of 98.42%, 89.83%, 88.79%, and 78.83%,, respectively. In Appendix B, I compare these metrics to those from similar recently published models, showing that my results are competitive with, or exceed, their performance.[30]

While each natural language model performs well, the geoparsing and data cleaning steps also influence the model's overall accuracy. Since these steps do not involve trained models, they cannot be evaluated using confusion matrices. Although I cannot calculate recall for the complete process, I can estimate precision for the non-imputed: the proportion of cases where the model predicted a criminal enterprise to be present in a given municipality and year, and that prediction was actually correct. To do this, I randomly sample 250 cases where the model predicted criminal presence, review the associated articles, and assess whether the prediction is supported by the text. If at least one article confirms the prediction, I label it as a true positive. Using this approach, I calculate a precision score of 71.2%.

External Validation

To externally validate my model, I compare my data to the Mapping Criminal Organizations (MCO) project, which tracks the presence of 32 criminal enterprises (CEs) between 2007 and 2015 at the state-month level. The MCO project draws from 60 documents from 11 expert and government sources (Signoret et al., 2021), making it particularly useful for external validation because it is not solely based on newspaper data, thus helping to mitigate concerns about media coverage biases.

However, the MCO data has certain limitations. First, its geographic coverage is broader, at the state level. Since my data is at the CE-municipality-year level, I convert both panels to the CE-state-year level for comparison. The panel is also unbalanced across CEs, as not all sources cover the same organizations or time periods. To ensure a robust comparison, I restrict the analysis to state-CE-year combinations where multiple sources confirm the presence of an organization. This allows me to exclude observations where a CE was classified

[29] The precision, recall, and F1 score of each sub-model is computed using the confusion matrices in Appendix Figures C2, C3 and C4.

[30] In this context, "similar" refers to models used for comparable tasks, such as predicting the topic of newspaper articles in non-English languages for text classification, or identifying entities like locations in large text corpora for named entity recognition (NER).

as having only a minor presence.[31]

While the MCO data is valuable for externally validating the newspaper data, the news data provides a more comprehensive picture of CE operations. It tracks more organizations at the municipality level and spans a longer time frame. All but three of the organizations tracked in the MCO data are also present in my data. However, over the same period, I track 375 groups that are not tracked in the MCO data.

The confusion matrix in Figure C5 provides insight into the external validity of my model. In terms of recall, the model performs well, with a score of 75.0%, correctly identifying a large proportion of the year-CE-state tuples tracked by the MCO data. Recall is particularly important in this context, as my primary concern is whether the model is missing organizations that should be present in specific states and years, potentially due to media suppression or uneven coverage across urban and rural areas. A relatively low recall would suggest that the newspaper data fails to capture key instances of CE activity—something we are especially worried about given the capacity for organizations to repress coverage through violence.

Precision, on the other hand, is less informative in this setting. It is more reasonable to assume that the MCO data accurately tracks CEs where it claims they are present, but is incomplete in terms of tracking all relevant instances of CE presence. Nonetheless, my model demonstrates strong precision of 69.9%, with relatively few false positives, which suggests that it is not systematically over-predicting the presence of criminal organizations.

Overall, the F1 score of 72.4%, which balances precision and recall, indicates that my model provides a robust measure of CE presence. Despite the challenges with the precision metric, the model performs well against the MCO benchmark, providing confidence in its ability to track criminal enterprises comprehensively over time and across municipalities. This compares favourably to the external validation of other models.[32] However, some concerns remain. Since I rely on a synthetic difference-in-differences (SDID) approach in my empirical strategy, potential errors in capturing criminal presence—particularly underreporting across both regions and years—could influence the assumption of parallel trends between treated and synthetic control municipalities. If these reporting errors vary systematically over time and across different areas, they may introduce bias in estimating the effects of key events, such as kingpin captures. Thus, after introducing the empirical strategy, I con-

[31] Signoret et al. (2021) defines a group's presence as minor if the CE entered the territory occasionally, without being based there or engaging in recurring operations, or if the group had operations only in a small part of the territory or in a limited segment of illegal markets.

[32] Other panels on CE presence do not report these external validation metrics. However, Sobrino (2019) reports accuracy using data from the Drug Enforcement Administration as the ground truth. She reports an average annual accuracy of 44%. The accuracy of my model when compared to the MCO data is 60%

duct a series of robustness checks in Section 5.3 to assess the extent to which these potential biases affect my empirical results.

3.2 Homicides Data

To estimate the effect of high-profile captures and deaths on violence, I use monthly data on homicides from the Mexican statistical agency (INEGI) between 2005 and 2019. This data is derived from death certificates, and is regarded as more trustworthy than the homicide rates published in the National Public Security System, which come from the number of police reports for homicide. Valle-Jones (2011) shows that several known massacres are missing from the police data, a problem that is far less severe in the INEGI data.[33] I focus on homicides because unlike other types of crime which are under-reported (e.g. kidnapping), observing a homicide does not depend on victims reporting the crime to police. I convert the data into per capita terms using population data from the Mexican census that is linearly imputed between census years.

Figure 1 shows the rate of homicides per 100,000 population on a monthly basis since the beginning of 1990. Homicide rates decline prior to 2008, when they began to rise above historic levels, increasing 162% between 2008 and 2019. Because homicides began to rise rapidly shortly following Caledròn's election, and particularly after the first kingpin capture during his term occurred in 2008, many academics and journalists argue that increased enforcement against DTOs, and specifically the policy of high-value targeting, not only failed to meet its goals of reducing violence, but actually increased rates of violence by creating power vacuums which incentivize fighting (Ríos (2013), Calderón et al. (2015), Osorio (2015), Phillips (2015)).

One limitation of the homicide data is that it does not distinguish between homicides that are related to organized crime and those that are not. To demonstrate that the overall homicide rate is a good proxy for DTO-related homicides, I compare this data with shorter panel data on drug-related homicides from the National Council of Public Security.[34] In this dataset, a homicide is classified as DTO-related if one civilian killed another civilian and at least one of the parties was involved in the drug trade. I find that drug-related homicides accounted for 45.4% of overall homicides during the period it was tracked. While these homicides explained less than half of the overall total, they explain 96.5% of the temporal variation in homicide rates. This is demonstrated in Appendix Figure A2 which plots the

[33] Valle-Jones (2011) shows there are cases where the total number of recorded homicides in a state is less than the total number of victims of a massacre.

[34] I do not use this data for the main analysis because it only covers the period between January 2006 and December 2010, while the majority of my DTO presence data covers the period after.

time series of both homicide rates. Overall homicides follow the same overall trend as DTO-related homicides.

3.3 Kingpin Captures/Deaths

To identify a list of Mexican kingpins, I start with the list of all Mexicans sanctioned under the Foreign Narcotics Kingpin Designation Act. Passed in 2000, the Act gives the U.S. Presidency the power to identify significant foreign narcotics traffickers and their associated entities (e.g. companies and operatives), to deny them access to the United States' financial system, and to prohibit all trade and transactions between these entities and U.S. companies and individuals (National Archives and Records Administration, 2009).[35] The Kingpin Act Sanctions List is ideal because it is the only list of CE leaders that covers the same period as my data on CEs' areas of operations. One limitation of the data is that it only includes CE leaders who pose a threat to the United States (e.g. through international drug trafficking). While there are some Mexican CEs that do not meet this criteria, they are not the focus of HVT. For example, every kingpin included on the Mexican government's 2009 most wanted list, which was distributed to Mexican law enforcement, also belonged to a CE that was sanctioned under the Kingpin Act (Lessing, 2017). Furthermore, there is no database of kingpins that includes the names of leaders of these organizations.

I match these names to information about their capture or death (if it has occurred).[36] This includes the cause of death, and the year, quarter, and municipality of capture or death. After performing research on each name on the sanctions list using press releases released by the Office of Foreign Asset Control, I determine whether or not each person is a kingpin of a CE, and restrict my analysis to these individuals.[37] Kingpins are leaders at the highest rank of an organization and are described in the press as the leader or one of the leaders of the CE. I focus on these leaders (rather than include lower-ranked lieutenants) While there is no complete list of either kingpin or lieutenant captures, evidence suggests that the kingpin list is more complete. For example, 68% of the kingpins listed on Mexico's most wanted list

[35] The Act applies to kingpins of criminal organizations with an international presence, though Mexico leads in terms of designations made under the act by a wide margin (with 62 designations made between 2000 and 2019) (U.S. Government Accountability Office, 2019). The penalties for U.S.-based entities that violate the Kingpin Act are steep: Individuals who violate the Kingpin Act are subject to criminal penalties of up to 10 years in prison and/or fines pursuant to Title 18 of the U.S. Code. Entities that violate the Act face criminal penalties in the form of fines up to $10 million; officers, directors, or agents of an entity who knowingly participate in a violation of the Kingpin Act are subject to criminal penalties of up to 30 years in imprison and/or a $5 million fine. The Kingpin Act also provides for civil penalties of up to $1.075 million against individuals or entities that violate its provisions (National Archives and Records Administration, 2009)

[36] I use a combination of queries of the ProQuest database and Google News to acquire this information

[37] If news sources do not identify the municipality or date of a leader's capture, or if their capture occurs outside of Mexico, their name is removed. I also exclude individuals tied to groups not tracked in my data.

were also sanctioned under the Foreign Narcotics Kingpin Designation Act. For lieutenants, this number was just 37%.

Figure 2 shows the number of kingpin captures per year between 2000 and 2019. In total, I observe 18 municipalities with kingpin captures or deaths between 2008 and 2016. Two of these captures occurred in the same municipality, requiring me to drop the later from my sample, as SDID cannot accommodate multiple treatments per municipality. Most years have 2 or 3 kingpin captures, and are frequent both during the beginning and end of the panel. This demonstrates that even though HVT is primarily associated with President Felipe Caledrón, subsequent administrations have continued the practice. Figure 4 shows the number of municipalities with a kingpin capture or death over the sample period on a state level. Many of the captures and deaths occur in areas that are strategic for drug trafficking, such as states on the U.S.-Mexico border or the highly fertile states of Sinaloa, Michoacán, and Jalisco, which are suitable for growing drug crops (poppies and cannabis).

By focusing on kingpin captures, I am able to gather more extensive information on the circumstances leading to their capture, as reporting about these individuals is more detailed. I use this information to select the subset of kingpin captures where treatment is plausibly exogenous. In particular, the occurrence and timing of a kingpin capture or death in a municipality should not be driven by changes in homicide rates, market structure or other local characteristics that might influence both the likelihood of a capture attempt and the outcomes of interest. Some potential violations include the possibility that the capture attempt was anticipated by the affected CE or their rivals, which could affect their market entry/exit decisions or their decision to engage in turf war. Additionally, if capture attempts were made in response to increasing violence in the region in which they occurred, I would be unable to leverage the differential timing of the events to estimate the causal effects of kingpin captures/deaths.

To address this, I report the details of each capture, as reported by news media, in Appendix Table A2. While the level of detail in the reporting varies across captures, the majority of the cases for which details exist support Lessing (2017)'s claim that in Mexico, HVT was by and large, exogenous. In many of these cases, the kingpin was captured or killed shortly after authorities obtained the required intelligence needed to locate them. The majority of capture attempts concluded quickly with either the successful capture or death of the kingpin. Some occurred in restaurants or other public settings, suggesting that the arrest or attempted arrest was unknown to the kingpin in advance.

However, I remove five kingpin captures from the final sample due to endogeneity concerns. Specifically, I exclude captures which meet the following criteria:

1. Prior to the kingpin's capture/death, capture attempts had been made more than one month prior to the event, but they failed (Antonio Cárdenas Guillén, Joaquín Archivaldo Guzmán Loera).

2. Government sources had falsely claimed that the kingpin died prior to their actual death (Nazario Moreno González).

3. The kingpin's capture occured during Caledrón's administration in states affected by Operation Northern Lynx after the beginning of the operation. There is evidence that this operation occured in response to civilian massacres by Los Zetas (Jorge Eduardo Costilla Sánchez, Iván Velásquez Caballero).[38]

After removing these kingpins, I am left with a sample of 12 kingpin captures and deaths. The leader, location, and timing of these events is reported in Table 1. To show that the criteria are not chosen selectively to generate the results, I include these events as a robustness check in Section 5.3.4.

To further support the argument that treatment is exogenous, I plot the monthly homicide rate per 100,000 population in the municipality of capture or death in the 24 months before and after each event for each of the remaining 12 kingpin captures in Appendix Figure A1. If kingpin captures and deaths were positively selected on the basis of rising violence in the municipality prior to the capture, we would expect that homicides would be rising in the months immediately prior to the capture or death. The figure suggests this is not the case, and there are many examples in which the homicide rate was falling in the months prior to capture or death, before sharply rising following the event.

3.4 Data For Control Variables

During the observation window (2005-2019), several large shocks to drug markets changed the relative profitability of controlling territory in across Mexican municipalities. In 2010, the U.S. Food and Drug Administration approved a new abuse-deterrent formulation of OxyContin, a drug largely responsible for fuelling the first wave of the opioid epidemic (Arteaga and Barone, 2022). As a result, U.S. demand for heroin increased, which increased the profitability of trafficking heroin for Mexican CEs and incentivized turf war (Sobrino, 2019). Heroin overdoses in the United States (the primary market for Mexican CEs) peaked in 2017 and were quickly surpassed by synthetic drug overdoses, especially those involving

[38] See Section 2.3 for a discussion of this operation and the political context surrounding it.

fentanyl and fentanyl-analogues (Ciccarone, 2021).[39] Mexican CEs produce these drugs in illicit laboratories using chemical precursors imported from Asia through Pacific ports (National Drug Intelligence Center, 2000). The growing market for synthetic drugs has increased the value of territorial control in municipalities with Pacific ports. Finally, as U.S. states began to legalize cannabis, Mexican CEs trafficked much less of the drug over the southern border, reflecting lower profitability due to the presence of a legal substitute (Mateos Zúñiga and Shirk, 2022).

If the timing and location of these shocks is correlated with that of kingpin captures, then my estimates of the effect of HVT on market structure and homicide rates will be biased. To ensure my results are not driven by these shocks, I interact time indicators with a measure of opium suitability, cannabis suitability, and an indicator of whether the municipality contains a Pacific port. I construct suitability indicies for opium and cannabis in lieu of using data on cultivation directly because the decision on where to cultivate is plausibly influenced by the strength of law enforcement in these areas. Instead, for both opium and cannabis, I regress plantation data on the average suitability of growing 312 (non-drug) crops from the Food and Agriculture Organization ECOCROP database and use the fitted value as an exogenous measure of suitability.[40]

Formally, the suitability index for growing opium poppy (and an analogous measure for cannabis) in municipality m is the fitted value from the following regression equation:

$$arcsinh(poppy\_plantations_m) = \alpha + \sum_{p=1}^{P} \beta_p \overline{suit}_{pm} + e_m$$

Where $\overline{suit}_{pm}$ is the predicted average output of crop p in municipality m, and $poppy\_plantations_m$ is the total number of illicit poppy plantations detected by authorities in municipality m between 1990 and 2020.

3.5 Restrictions on Control Municipalities

I restrict the set of control municipalities to exclude those that present threats to identification. First, I exclude municipalities in which captures or deaths occurred within the observation window but outside of the event window so that these capture events do not contaminate the results. For the same reason, I also exclude municipalities in which captures

[39] The sudden prominence of fentanyl in illicit drug markets can be explained through a confluence of supply-side factors, including more efficient synthesis methods and the rise of internet communication and commerce (Pardo et al., 2019)

[40] Specifically, I use data on the number of plantations from SEDENA (Secretary of National Defense), SEMAR (Mexican Navy), GN (National Guard) and the PF (Federal Police).

were endogenous (discussed in Section 3.3) from the set of possible controls. Additionally, it is possible that a leader's capture or death not only affects market structure and violence in the municipality of capture or death, but also in those that are geographically proximate. If this is the case, the stable unit treatment assumption, which requires that the response of a particular unit depends only on the treatment to which it was assigned, would be violated, generating biased estimates. To address this, I drop all municipalities bordering any treated municipality and all of the municipalities that border the bordering municipalities from the sample of eligible control municipalities to ensure that municipalities that are likely to be affected by spillovers are not included in the synthetic control. I do the same for the potentially endogenous capture municipalities. Mexico's 2457 municipalities are divided as follows: 12 are in the treatment group, 2151 are in the control group, and 294 are in neither group and are thus excluded from the analysis.

3.6 Stylized Facts

My dataset allows me to measure the presence of criminal enterprises (CEs) in Mexico in a way that previous datasets have not. Using natural language processing techniques, I track the activities of 408 unique criminal enterprises at the municipality-year level between 2005 and 2019. This approach is particularly effective in identifying smaller, less prominent organizations that have become increasingly common in Mexico's criminal landscape. Unlike previous studies, which often rely on pre-defined lists or focus on a limited number of larger organizations, my data provides rich information on each criminal enterprise, their geographic operations, and their relationships with other groups.

This rich dataset enables me to uncover new stylized facts surrounding the nature of fragmentation among criminal enterprises in Mexico. By tracking a greater number of organizations, including smaller and less prominent groups, I capture dynamics that were previously difficult to observe, particularly the ways in which criminal enterprises compete, fragment, and reorganize. These insights provide a more comprehensive understanding of how market structure evolves in response to external pressures, such as kingpin capture and death. In the following sections, I will test whether these observed trends are driven by these leadership disruptions and explore their broader effects on violence and control in affected regions.

The insights from the stylized facts reveal that market concentration within Mexican municipalities has been falling over time, driven primarily by intensified competition among criminal enterprises already operating in key areas, rather than geographic expansion. Additionally, the growth in the number of criminal enterprises is largely attributed to the

splintering of existing organizations.

1. Within Mexican municipalities, market concentration of criminal enterprises has fallen over time.

A key finding in this study is the evolving market concentration of criminal organizations across Mexican municipalities. Previous analyses, limited by their inability to capture smaller criminal entities, have largely focused on the dominance of major cartels. This approach overlooked the increasingly significant role of splinter groups, which now shape Mexico's criminal landscape. I use the Herfindahl-Hirschman Index (HHI) to capture these market dynamics, measuring concentration by summing the squared market shares of each criminal organization present within a municipality. Importantly, my analysis is the first to apply the HHI to Mexico's criminal market on a national scale.[41] The HHI for municipality m in year t is calculated as:

$$\text{HHI}_{mt} = \sum_{g} s_{gmt}^2 \tag{1}$$

where s_{gmt} is the market share of criminal organization g in municipality m during year t. A higher HHI indicates a more concentrated market dominated by a few organizations, while a lower HHI reflects greater competition among multiple groups. In the absence of specific data on revenues or criminal output, I assume equal market shares for all present organizations.[42]

Figure 5 shows that from 2005 to 2019, market concentration notably declined, with the HHI dropping from around 0.90 to 0.75, reflecting a transition from a moderately concentrated market toward greater competition. According to U.S. Department of Justice (2023), markets with an HHI in this range (0.70 to 0.90) are still considered moderately concentrated, meaning that while competition increased, the criminal landscape was still dominated by a relatively small number of groups during this period. This trend coincides with the intensified efforts to capture kingpins, as displayed in Figure 2. After 2016, the decline in market concentration levels off, with the estimates falling within the 95% confidence interval, suggesting a stabilization of market dynamics. Whether this stabilization is directly tied to kingpin captures is an empirical question that I address in Section 5, using an event study

[41] While Sobrino (2019) calculates the HHI in opium-growing municipalities, this broader application to the entire criminal market is novel.

[42] A limitation of my data is the inability to observe inter-organizational alliances. Thus, all present organizations are assumed to compete, and any newly independent splinter group is treated as a competitor only after its split from its parent organization.

approach to estimate the causal effects.

2. The observed shift is primarily driven by growing competition among criminal organizations within municipalities where they were already active, rather than by expansion into new areas.

Understanding whether the increase in the number of criminal enterprises is driven by intensified activity within municipalities or by expansion into new areas is crucial for identifying the mechanisms behind this growth. To analyze this, I decompose the overall change into two parts: the intensive margin and the extensive margin. This helps distinguish between growth in areas where criminal enterprises were already present and growth in new areas.

I divide the data into two periods: 2005–2012 and 2013–2019, corresponding to the first and second halves of the observation window. Appendix Figure A3 visually illustrates these shifts across Mexico's municipalities, showing the number of criminal enterprises operating in each region during both periods. The maps show that while some areas saw new entrants, many municipalities with criminal enterprises in the early period continued to experience competition later on, suggesting that growth is primarily concentrated within existing areas of operation. Most of the growth is driven by the intensive margin, accounting for 65.5% of the total change.[43] This reflects the increasing number of criminal enterprises operating within municipalities that were already active in the earlier period, rather than a widespread geographic expansion into new areas.

This trend suggests a growing concentration of criminal organizations in key regions, likely reflecting intensified competition over valuable resources, such as smuggling routes and other strategic assets. As new entrants increasingly operate in municipalities where criminal enterprises were already present, competition has likely escalated within these areas. The result is a crowded criminal landscape, where increased activity is driven more by the concentration of enterprises in established regions than by expansion into untouched areas. By examining the causes and consequences of kingpin captures and deaths in Section 5, I will use event study estimates to demonstrate that HVT was responsible for some of this change.

[43] The changes in the intensive and extensive margins were calculated by splitting municipalities into those where criminal enterprises were present in both periods (intensive margin) and those where they entered or exited (extensive margin). For each group, the total change in the number of criminal enterprises was calculated by summing over municipalities, and the contribution to the total change was expressed as a percentage. The combined contributions of the intensive and extensive margins sum to 100%.

3. Falling market concentration within municipalities is driven by splintering criminal organizations.

The growth in the number of criminal enterprises operating in Mexican municipalities over time has been largely driven by the splintering of existing organizations. Figure 6 shows the average number of criminal enterprises by type across years: those that were always cartels, those that were always independent, and those that splintered from larger organizations. The data span from 2005 to 2019, providing a clear view of how these groups evolved over time.

The figure demonstrates that splinter groups have experienced the most significant growth, particularly since 2006, after the launch of intensified kingpin targeting strategies under the Calderón administration. In contrast, the number of criminal organizations that were always independent or always operated as cartels has remained relatively stable, with only modest increases. This makes it evident that the sharp rise in the number of criminal enterprises is primarily driven by the fragmentation of larger organizations into splinter groups, rather than by the formation of entirely new or independent organizations.

I identify these splinter groups by tracking when a previously unified organization breaks apart, using newspaper data to map out the year in which splintering occurred. By distinguishing between criminal enterprises that have always been independent, those that consistently acted as cartels, and those that emerged from splintering, I can assess how fragmentation has contributed to the growth of the number of CEs across Mexican municipalities.

Because my estimates rely on newspaper reporting to detect the presence of criminal enterprises, one concern is that increases in coverage across time and space could drive the changes discussed above. In particular, even if market structure remained stable, more intensive press coverage in certain places or periods would make it more likely that criminal groups are detected. To help separate true changes in market structure from changes in detection, I use the number of non–organized crime articles mentioning each municipality in each year as a proxy for general newspaper coverage. These articles are less likely to be influenced by organized crime itself, making them a useful way to adjust for shifts in media attention that are unrelated to the actual presence of criminal enterprises. To derive this data, I start with the set of articles that the text classification model predicts are not about organized crime. Then I apply the geoparsing algorithm to these data and use the results to build the panel.[44]

[44]The text classification model is discussed in detail in Appendix Section B.1 and the geoparsing algorithm is discussed in Appendix Section B.4.

To adjust for fluctuations in newspaper coverage, I estimate the following regression:

$$\mathrm{y}_{mt} = \alpha + \beta \cdot \mathrm{X}_{mt} + \gamma_m + \epsilon_{mt} \tag{2}$$

where y_{mt} is the number of detected criminal enterprises (or in the case Figure 6 the number of detected criminal enterprises of a particular type) in municipality m and year t, and X_{mt} is the number of non–organized crime newspaper articles mentioning municipality m in year t. The term γ_m denotes municipality fixed effects. This specification captures the average within-municipality relationship between newspaper coverage and observed criminal presence over time. I then subtract the predicted contribution of article volume—relative to each municipality's average coverage level—to isolate residual variation in detected criminal presence.[45]

To implement this adjustment, I first calculate the predicted number of detected criminal enterprises from equation (1) using the estimated coefficient $\hat{\beta}$ and each municipality's deviation in article volume from its own average. Specifically, I compute:

$$\tilde{y}_{mt} = \mathrm{y}_{mt} - \hat{\beta} \cdot (\mathrm{X}_{mt} - \bar{\mathrm{X}}_m) \tag{3}$$

where $\bar{\mathrm{X}}_m$ is the average number of articles mentioning municipality m across all years. This step removes the portion of observed variation that can be explained by temporary increases or decreases in coverage, relative to that municipality's own average.

Finally, I apply a cleaning rule to the adjusted CE counts to ensure they are compatible with concentration metrics like the HHI and do not produce negative CE counts.[46] Specifically, I obtain the adjust CE counts, which I use in my $\hat{y}_{mt}$, which I use in my adjusted stylized facts, in the following manner:

$$\hat{y}_{mt} = \begin{cases} 0 & \text{if } \tilde{y}_{mt} < 0.5 \\ 1 & \text{if } 0.5 \leq \tilde{y}_{mt} < 1 \\ \tilde{y}_{mt} & \text{if } \tilde{y}_{mt} \geq 1 \end{cases}$$

The adjusted stylized fact graphs are plotted in Appendix Figures A4-A6. The results

[45] I exclude year fixed effects because they would absorb national-level time trends in both coverage and detected criminal activity. Since my goal is to assess how market structure changes over time, including year fixed effects would strip out meaningful temporal variation in criminal presence and result in over-correction.

[46] This rule addresses two concerns. First, the adjustment procedure can yield negative values in municipalities with low baseline coverage or large increases in press volume. Second, concentration indices like the HHI require a nonzero count of competitors. Additionally, values below 1 would lead to undefined or misleading results and division by the decimal could result in HHIs that are above the maximum value of 1. The rule preserves meaningful variation while ensuring that the adjusted counts remain interpretable and usable.

graphs depict the same overall trends, suggesting that the stylized facts reflect meaningful changes in criminal presence over time, rather than being driven by growth in newspaper coverage that could result in more CEs being detected.[47] The largest change is to Figure 5, where the decline in the HHI is not as large in early years in the adjusted figure (Appendix Figure A4) due to a lower article count during this time period.[48] However, HHI still declines over time and as Appendix Table A13 shows, the difference between mean HHI in 2019 and 2005 is statistically significant.

4 Empirical Method

4.1 Empirical Strategy

My goal is to identify the causal effect of HVT on the number of criminal enterprises operating in the municipality of capture or death and the policy's subsequent effects on homicide rates. Specifically, I aim to recover an estimate for the average treatment effect on the treated: I estimate how market structure and violence are changing in municipalities that had a kingpin capture or death, relative to a counterfactual in which it did not occur.

To do so, I follow Arkhangelsky et al. (2021) and implement the synthetic difference-in-differences estimator.[49] SDID combines desirable features of both synthetic control methods and traditional two-way fixed effects models. Like synthetic control, SDID constructs a weighted synthetic control group by optimally choosing unit weights, aligning treated and control units' trends in the pre-treatment period for each treated cohort. However, unlike synthetic control, it does not minimize average differences in levels in the pre-treatment period. Instead, SDID selects weights across control units to approximate parallel trends in the pre-treatment period, aiming to match treated units' trends as closely as possible, assuming suitable controls are available. Like DID, the method is invariant to additive unit-level shifts, facilitates the inclusion of time-varying controls, and allows for large-panel inference. SDID also estimates a set of time weights that optimally reweight the pre-treatment periods, assigning lower weights to periods that are unusual for the control municipalities relative to the post-treatment period.

SDID is useful in contexts like mine where treated units are substantially different from

[47] The decomposition of growth along the intensive and extensive margins that accompanies Figure A3 is very similar, with 63.4% of the variation being due to growth along the intensive margin after adjustment.

[48] The adjustment has a much larger effect for HHI than it does for the other figures due to the non-linearity of HHI.

[49] Specifically, I follow Clarke et al. (2023), who demonstrates how to estimate Arkhangelsky et al. (2021) in a setting with staggered treatment, and Ciccia (2024), who disaggregates these estimates into dynamic treatment effects and develops the Stata package *sdid_event* to implement the procedure.

potential controls. Table 2 reports separate summary statistics for municipalities with kingpin captures or deaths in the pre-treatment period, as compared to the entire sample. The summary statistics indicate substantial disparities in variables such as the number of criminal enterprises, population, and homicide rates per 100,000 between treated units and all municipalities. These differences suggest that treated units may follow different baseline trends and characteristics compared to the broader sample, which could violate the parallel trends assumption required by TWFE. SDID, by constructing a synthetic control for each treated unit and adjusting for staggered treatment timing, allows for better alignment with the characteristics of treated units.

The core identification assumption in SDID is that for each treatment date, we can construct a synthetic control that either (a) follows a similar outcome trend to treated units in the absence of treatment (correctly specified outcome model) *or* (b) accurately reflects the factors influencing treatment assignment, thereby isolating the treatment effect from other factors. This framework allows SDID to attribute post-treatment differences in outcomes between treated units and their synthetic controls to the treatment effect, provided that either we can create a synthetic control that represents counterfactual outcomes or the treatment assignment model is well-specified. Since counterfactual outcomes for municipalities without kingpin captures or deaths are unobserved, neither the outcome model nor treatment model assumptions are directly testable. In Section 3.3, I provide descriptive evidence to support the validity of the estimates by detailing the circumstances leading up to each capture attempt, helping to address concerns about positive selection. Additionally, event study estimates (Equation 10) facilitate a more detailed visual test for pre-trends that accounts for covariates, offering indirect support for the outcome model's parallel trends assumption.

4.2 Empirical Specification

Obtaining SDID estimates in a staggered context requires estimating separate treatment effects $\hat{\tau}^{sdid,e}$ for each cohort $e \in \{e_1, ..., E\}$. Each cohort includes N_{co} never-treated control municipalities and N_{tr}^{e} treated municipalities that had their first kingpin capture at time $t = e$, with $N^e = N_{tr}^{e} + N_{co}$. Furthermore, I let T_{pre}, T_{post} and T denote the number of pre, post, and total number of periods, respectively. For market structure related outcomes, t indexes years and for homicide rates, t indexes quarters. I let $l \in \{\underline{j}, \bar{j}\}$ index relative time and restrict the data to a an event window of $\underline{j}$ periods before and $\bar{j}$ periods after treatment such that both quarterly and annual outcomes cover a post-treatment period of 4 years.[50] The cohort-specific ATT estimates, $\hat{\tau}^{sdid,e}$ are chosen to solve the following optimization

[50] I set $\underline{j} = 3$ and $\bar{j} = 4$ for market structure and $\underline{j} = 5$ and $\bar{j} = 16$ for homicide rates.

problem:

$$\underset{\tau^{sdid,e},\mu^e,\alpha^e,\beta^e,\delta^e}{\arg\min} \sum_{m=1}^{M}\sum_{t=1}^{T^e} (Y_{mt}^e - \mu^e - \alpha_m^e - \beta_t^e - X_{mt}^e\delta^e - D_{mt}^e\tau^e)^2 \,\hat{\omega}_m^{sdid,e}\hat{\lambda}_t^{sdid,e} \tag{4}$$

where Y is the outcome of interest, X is a vector of time-varying controls, and D indicates whether municipality m had its first kingpin capture by time t. The optimization problem is similar to that of TWFE, with the exception of the weights.[51] $\hat{\omega}_m^{sdid,e}$ are the unit weights that minimize the average squared difference in trends in between treated and control groups in the pre-treatment period for cohort e, subject to a regularization parameter. The regularization parameter prevents overfitting by increasing the dispersion of the weights to ensure their uniqueness. $\hat{\lambda}_t^{sdid,e}$ are the time weights chosen to minimize the sum of squared differences between the time-weighted average pre-treatment outcome in control municipalities and their unweighted average outcomes in the post-treatment periods. The estimation of these weights is defined in detail in Appendix Section C.1.

The resulting estimator is a cohort-specific double-difference estimator that incorporates the SDID weights. $\hat{\tau}^{sdid,e}$ can be written as the difference in adjusted outcomes between the average treated municipality in cohort e and a synthetic control unit that averages the never-treated controls using an optimal set of weights:

$$\hat{\tau}^{sdid,e} = \frac{1}{N_{tr}^e}\sum_{m=N_{co}+1}^{N^e}\hat{\gamma}_m^e - \sum_{m=1}^{N_{co}}\hat{\omega}_m^{sdid,e}\hat{\gamma}_m^e \tag{5}$$

where

$$\hat{\gamma}_m^e = \frac{1}{T}\sum_{t=T_{pre}+1}^{T}\tilde{Y}_{mt}^e - \sum_{t=1}^{T_{pre}}\hat{\lambda}_t^{sdid,e}\tilde{Y}_{mt}^e$$

and $\tilde{Y}_{mt}^e$ are the residuals from from the regression of $\tilde{Y}_{mt}^e$ on the vector of time-variant covariates X_{mt}^e. I aggregate the cohort-specific estimates to estimate the cohort average treatment effect on the treated for outcome Y by weighting each cohort's treatment effect estimate by its share of the total number of post-treatment periods for treated municipalities:

$$\hat{\tau}^{sdid} = \sum_{e=e_1}^{E}\theta^e \times \hat{\tau}^{sdid,e} \tag{6}$$

where

[51] In contrast, TWFE uses uniform weights.

$$\theta^e = \frac{N^e_{tr} T_{post}}{\sum_{e=e_1}^{E} N^e_{tr} T_{post}} \tag{7}$$

With these results in hand, I estimate the 4 year average percentage change in outcomes for treated municipalities:

$$\%\Delta Y = \frac{\hat{\tau}^{sdid}}{\bar{Y}_{pre}} * 100 \tag{8}$$

where $\bar{Y}_{pre}$ denotes the $\hat{\lambda}^{sdid}_t \times \theta^e$ weighted average value of Y_{mt} in the pre-treatment period for treated municipalities.

To test for selection bias, I allow trends in adjusted outcomes $\tilde{Y}_{mt}$ to deviate between the treatment and control group in the pre-treatment period within each cohort and disaggregate the cohort-specific treatment effects into dynamic estimates. I estimate the treatment effect for municipalities in cohort e $l \in \{\underline{j}, \bar{j}\}$ periods before and after a kingpin capture or death:

$$\hat{\tau}^{sdid,e}_l = \frac{1}{N^{e+l}_{tr}} \sum_{m=N^l_{co}+1}^{N^{e+l}} \hat{\gamma}^{e+l}_m - \sum_{m=1}^{N^l_{co}} \hat{\omega}^{sdid,e}_m \hat{\gamma}^{e+l}_m \tag{9}$$

where

$$\hat{\gamma}^{e+l}_m = \tilde{Y}^{e+l}_{mt} - \sum_{t=1}^{T_{pre}} \hat{\lambda}^{sdid,e}_t \tilde{Y}^{e}_{mt}$$

I then retrieve cohort average dynamic treatment effects using weights θ^e:

$$\hat{\tau}^{sdid}_l = \sum_{e=e_1}^{E} \theta^e \times \hat{\tau}^{sdid,e}_l \tag{10}$$

This approach varies from a standard event study specification in which treatment effects are estimated relative to an arbitrary baseline (often one year prior to treatment). Furthermore, since the weights are selected using the pre-period outcomes, the SDID weights, and thus the synthetic control, is different for each outcome.[52] In the case of SDID, coefficients are computed for each relative time period in the event window relative to their weighted pre-period average. I calculate standard errors for all estimates using the block bootstrap variance estimation algorithm described by Clarke et al. (2023) with 100 replications and state level clustering to account for correlation of treatment status within states.

[52] This is important to consider when I explore heterogeneity in outcomes, as the individual treatment effects do not sum up to the total effect in these cases.

5 Results

5.1 Market Entry and Exit

I begin by estimating the average effect of a kingpin capture or death on the market structure of organized crime in treated municipalities. Figure 7 shows the event study estimates for the total number of criminal enterprises. The results show that the number of CEs trended similarly in treated municipalities and the weighted set of control municipalities in the 3 years prior to treatment. Afterwards, the number of CEs in treated municipalities increased, with the policy's effect diminishing over time. In Table 3, I show that on average, a kingpin capture or death causes an additional 1.15 CEs to operate in treated municipalities in the 4 years post-treatment: A 54.25% increase relative to the pre-treatment mean.

There are two potential mechanisms which could explain these results. First, HVT could incentivize rival groups to enter into affected municipalities. Osorio (2015), for example, proposes a model in which HVT weakens the organizational capacity of the targeted group. As a result, the costs associated with seizing the group's territory decline, encouraging competitors to engage in turf wars.[53] Second, kingpin captures and deaths may provoke inter-group conflict by creating power vacuums, which in turn lead to the formation of new factions or cause existing factions to switch affiliations and contest the group's territory. To assess these mechanisms, it is essential to distinguish between criminal enterprises (CEs) that have splintered and those that have not. Additionally, I must differentiate between recently splintered CEs and those that splintered further in the past, as the presence of a splinter group in a municipality many years after it first emerged (anywhere in Mexico) is more likely indicative of market entry than faction creation.

To do this, I define new splinters as either new factions of a criminal enterprise (CE) or former factions that recently split from a CE. Formally, a CE is classified as a new splinter of CE_A at time t if $t_A \leq t \leq t_A + 4$, and it meets one of the following criteria: (1) it was a former faction of CE_A that ended its affiliation with CE_A at time t_A, or (2) it was a faction of CE_A that first appeared in the data at time t_A. Restricting new splinters to the period $t_A \leq t \leq t_A + 4$ helps address the possibility that a former faction is entering to contest territory, rather than in response to intra-CE conflict. This is more likely if a significant amount of time has passed since the split or the group's initial appearance in the data. If neither of these conditions is met, the group is classified as a non-new splinter CE.

To test whether HVT induces entry of new groups, I separately estimate the policy's effect

[53] For example, the military resources required to successfully defeat the incumbent may decrease following the kingpin's capture or death.

on the number of non-new splinter CEs operating the municipality of treatment according to whether or not the group existed in the period prior to treatment. This helps distinguish between effects stemming from new entry, versus effects stemming from existing groups being less likely to leave their territory following treatment. While I cannot rule out the latter effect, the results in Figure A7 show the effect is primarily driven by new entry and concentrated in the year of treatment, though the results are persistent for several years following the change.

To test for splintering, I estimate the policy's effect on two separate outcomes, relative to synthetic control municipalities: (1) the number of locally-operating splinters of CEs that were present in each municipality in the year prior to treatment, and (2) the number of locally-operating splinters of CEs that were not previously present. This comparison allows me to assess whether the splintering of existing CEs occurs disproportionately in treated municipalities following treatment. If the number of splinters from CEs not previously present increases, it suggests that new groups are entering the municipality to contest the territory, rather than splintering from existing organizations. I report the event study estimates in Figure A8. There is no effect for unrelated new splinters. While the estimates for new splinters of existing groups are imprecise, and I cannot rule out these effects, their magnitude is much smaller than the effect for new entry, suggesting that kingpin captures primarily increase violence by increasing competition between existing criminal enterprises.

5.2 Homicides

Next, I explore the implications of the changes in market structure on homicide rates. If market entry is driven by territorial disputes, we should expect these results to be accompanied by a corresponding increase in homicide rates following executive capture or death. The event study estimates, reported in Figure 8, are consistent with this narrative. The results show that on average, in the five quarters preceding a kingpin capture or death, the trend in quarterly homicide rates is parallel. However following treatment, homicides sharply rise in treated municipalities. While the effect is largest in the quarter of capture or death, homicide rates remain elevated for 2 years following treatment before returning to baseline. In Table 4, I show that in the 4 years following a kingpin capture or death, treatment induces an additional 1.60 homicides per 100,000 population: A 31.78% increase over the pre-treatment mean. The results contribute to a large body of literature that demonstrates that competition between CEs is often violent in nature, in sharp contrast with legal industries (e.g. Levitt and Venkatesh (2000), Ríos (2013), Castillo et al. (2020), Sobrino (2019)). Scholars argue that competition is inherently unstable for illegal industries because

there are no formal dispute mechanisms to deal with conflict between organizations.

A key question is whether or not these deaths are concentrated amongst CE members. While the implications for civilian welfare are negative in both cases,[54] if turf-war related violence spills over to the civilian population, this is cause for greater concern. While data covering CE-related homicides is not available for the majority of the observation window,[55] the complete homicides panel reports the age and sex for the vast majority of homicide victims.[56] I follow Calderón et al. (2015) and proxy for CE-related homicides using the homicide rate of men aged 15-39.[57] I separately estimate the effect of kingpin captures and deaths on the homicide rate for likely CE members and the broader population. Figure A9 presents these findings, showing that the homicide rate rises for both groups, increasing 41.58% for males aged 15-39 and 41.23% for others in the 4 years following treatment. These results align with earlier work by Daniele et al. (2023), which shows that turf war drives outmigration from affected areas to the United States, highlighting its substantial toll on the civilian population.

5.3 Robustness Checks

5.3.1 Assessing Pre-Treatment Fit and the Credibility of the Counterfactual

While the main event-study results provide a clear summary of average treatment effects on the treated, they may conceal underlying issues. An aggregate event study with parallel pre-trends and post-treatment divergence can mask problems like a poor pre-treatment match for certain cohorts or an over-reliance on a few control units in constructing counterfactual outcomes.

To address these limitations, I also plot the raw outcomes over time for the average treated municipality in each cohort, alongside its synthetic control in Figures A10-A17. These cohort-specific graphs offer a granular view of how well each treated unit tracks its synthetic counterpart before treatment and where deviations occur afterward. The figures show that for most pre-treatment periods across cohorts and outcomes, the trends between

[54] For example, the violence has been shown to negatively affect the local economy in impacted municipalities (Utar, 2018).

[55] See Section 3.2 for a fuller discussion of this data and its limitations.

[56] For a small minority of homicides, either age or sex is omitted. These observations are dropped from the dataset, meaning that the observed estimates should be viewed as a lower bound of the observed effects. These missing observations account for 0.79% of the homicides observed in the panel, thus justifying this approach.

[57] Calderón et al. (2015) compares the count of homicides derived from death certificate data (as reported by the INEGI), which is also used in this paper, with data on drug trade related homicides from the National Council of Public Security. Through an exhaustive search of possible age and gender groupings and regression analyses, they determined that homicides for males aged 15-39 exhibited spatio-temporal variation that closely aligned with that of drug-related homicides.

treated municipalities and their matched synthetic controls are parallel prior to treatment. This demonstrates that the pre-treatment trends are not only parallel on aggregate (as shown in the event study graphs) but that generally, these parallel pre-treatment trends can be constructed for each outcome on a cohort-by-cohort basis, thus strengthening the credibility of the counterfactual outcomes. Additionally, there is no clear pattern surrounding which cohorts contribute the most to the results. While there is variation in the magnitude of treatment effects across cohorts and outcomes, it is clear from these cohort-specific figures that the effect is not driven by a small subset of treated municipalities.

An additional concern in synthetic control designs is weight dispersion. If most of the weight is assigned to a small subset of control units, any external shocks that impact those units during the treatment period could drive the observed differences, resulting in biased estimates. To address this concern, I plot histograms of the weights used for the synthetic counterfactual for each treatment date, along with a cumulative summation of the weights in Appendix Figures A18-A25. The graphs show that for most outcomes and cohorts, there is a high degree of dispersion of the weights, meaning that the synthetic control units are constructed by weighing the outcomes across many municipalities, with each being assigned a small weight. Additionally, I report the 10 largest weights for each cohort and outcome in Appendix Tables A3-A10. Across all cohorts for the main outcomes, the maximum weight assigned to any given municipality when constructing the counterfactual municipality is 34.46%. However, for most cohorts and outcomes, the most common weights are substantially smaller. Of the 64 cohorts across all main outcomes, there were only 13 municipalities with weights above 10% percent. Together, these results suggest that the synthetic counterfactuals are derived from a diversified set of municipalities.

5.3.2 Measurement Error in Newspaper Data

In Section 3.1.2, I demonstrated that, at the aggregate level, my newspaper-based data performs well when compared to data that is not based on newspaper coverage, as shown through the confusion matrix in Figure C5.However, what I have not yet examined are potential sources of non-classical measurement error that could exist in the newspaper data, which might not have been captured in the aggregate performance measures. For example, CEs might repress media coverage in areas they control, particularly after key events such as kingpin captures, leading to under-reporting of criminal enterprises in treated municipalities. Additionally, the geographic coverage of my newspaper database may change over time as I use all available articles in the database, with some years having more newspapers. If certain regions or periods have more comprehensive coverage, this could create false trends in the

data, potentially threatening identification by attributing changes in the reported number of criminal enterprises to treatment effects, when in reality they are driven by differences in reporting practices. In both cases, systematic underreporting could bias my results by either masking increases in criminal enterprise activity or exaggerating changes that are, in reality, artifacts of better reporting in certain years or regions.

Fixed effects help address some of these concerns by controlling for time-invariant factors that might influence reporting bias. Municipality fixed effects ensure that comparisons are made within the same municipalities over time, accounting for persistent differences in media presence or local conditions that might affect how criminal activities are reported across regions. Time fixed effects account for national shifts in reporting practices or media freedom that might affect all municipalities similarly over time. However, fixed effects do not address concerns related to time-varying or event-specific biases, such as media suppression by cartels following a kingpin capture or shifts in geographic coverage due to the inclusion of more newspapers in certain periods. These types of biases could still introduce systematic differences between the treated and control municipalities.

I first conduct robustness checks using a complementary dataset from the Mapping Criminal Organizations (MCO) Project, which tracks criminal organizations at the state level. While some of the MCO data is based on newspaper reports, I restrict my analysis to observations that are not identified using newspaper data alone. Although the MCO dataset lacks the fine-grained municipality-level detail, it provides a valuable benchmark to assess the robustness of my newspaper-based estimates.[58]

First, consider the possibility that media suppression may lead to non-classical measurement error in the data. Mexico is routinely ranked as one of the most dangerous countries in the world to be a journalist due to threats from organized crime, though the specific effects this has on the completeness and accuracy of reporting across locations are not fully known (RSF, 2024). Criminal organizations often target journalists reporting on cartel violence, drug trafficking, and government corruption, but their approach to controlling media narratives can be inconsistent (CPJ, 2010). For instance, criminal organizations frequently erect narcomantas—large banners displayed in prominent public spaces—to issue threats, assert dominance, or communicate with rivals, authorities, and the public. These banners, which openly acknowledge their presence, contrast with instances where cartels attempt to suppress media coverage, highlighting their selective engagement with the media.[59] This inconsistency presents a potential threat to my identification strategy if certain organizations—especially

[58] The MCO data is discussed in greater detail in Section 3.1.2, where it is introduced.

[59] For examples of narcomantas, see Appendix Figure A26.

those with greater resources—are able to systematically repress to hide their own activities. If prominent organizations consistently influence the media in this way, we would expect them to have systematically lower recall.

To address this concern, I conduct robustness checks by testing both the recall and precision of my newspaper-based data against the MCO data. Recall measures the extent to which the newspaper data accurately captures the presence of criminal enterprises, while precision evaluates whether reported organizations are actually present. By analyzing recall as across different levels of organizational prominence, I can assess whether more prominent organizations—those with the greatest resources to influence media coverage—are systematically under-reported.

To test this theory, I a proxy for the prominence of each criminal enterprise in the data. The measure is based on the proportion of observations in which the CE is reported across different states and years. Specifically, I calculate the prominence of each CE g in year t as:

$$P_{tg} = \frac{\sum_s \mathbb{1}_{st}^{\text{presence}}(g)}{\sum_{sg} \mathbb{1}_{st}^{\text{observed}}(g)} \tag{11}$$

where $\mathbb{1}_{st}^{\text{presence}}(g)$ is an indicator equal to 1 if CE g is indicated as present in state s at time t in the MCO data and $\mathbb{1}_{st}^{\text{observed}}(g)$ is an analogous measure for whether the group is tracked in that state at that time. This measure accounts for variations in data availability across states and years that are likely to affect recall. The overall prominence of each CE is then calculated by averaging its presence across all available years:

$$Prom_g = \frac{1}{T_g} \sum_t P_{tg}, \tag{12}$$

where T_g represents the total number of years in which CE g is observed. I categorize CEs into quintiles based on this measure and examine whether larger, more prominent organizations have systematically lower precision and recall. Appendix Figure A27 shows the mean recall rates for criminal enterprises (CEs) across quintiles of prominence.[60] One important observation is that recall is lower for the least prominent organizations (in the first quintile), highlighting a limitation of the data: smaller or less visible groups may be under-reported. However, as we move to more prominent organizations (quintiles 2-3), recall rates stabilize, and the differences between these groups become minimal. The most prominent organizations, those likely to be larger and more powerful, and thus have the greatest ability to suppress reporting, show the highest recall.

[60] The recall is weighted by the number of observations within each quintile.

A second potential source of bias stems from variation in the number of articles included in my dataset over time or inconsistencies in regional coverage. I address this by tracking the number of newspapers across years to ensure that changes in coverage are not correlated with treatment timing, and by comparing trends in my newspaper data to state-level data from the MCO Project. In Appendix Figure A28, I compute the recall of my data against the MCO data,[61] and observe that recall increases significantly in 2008, reflecting a large increase in the number of articles included in the database in that year (see Appendix Figure A29). However, after 2008, recall no longer mirrors the growth in articles or sources, suggesting that changes in coverage do not explain the trends in recall for most of the panel.

Fluctuations in recall across time are not inherently a threat to identification, as these can be absorbed by time fixed effects if they are common to all locations. However, the primary period of concern is pre-2008, as I only have one pre-2008 year of data from the MCO Project, making it difficult to fully assess whether trends were parallel prior to treatment. This limited availability of pre-treatment data complicates the identification strategy, as I cannot definitively capture pre-trends. A particular concern is the possibility that changes in media coverage occurred differentially between treated and control areas, which could bias my results. For example, my estimates could be inflated if media coverage increased disproportionately in treated areas or diminished in control areas, creating the appearance of a treatment effect driven by reporting differences rather than actual treatment.

In Appendix Figure A30, I assess recall trends by dividing states into those with above- and below-median populations based on their initial population levels in 2005. The figure shows that recall for both groups follows similar trajectories post-2008, with both groups experiencing an increase in recall between 2008 and 2011, followed by a decline. Importantly, while the trends are largely parallel after 2008, the lack of additional pre-treatment years makes it difficult to fully assess whether trends were truly parallel prior to treatment. This raises concerns that differences in media coverage prior to 2008 could affect identification, though the similar post-2008 recall patterns between the two groups mitigate this concern to some extent.

To further address this issue, I examine the sources added to the ProQuest database between 2008 and 2011—a critical period during which 16 new sources were added and 4 sources were removed. Importantly, all of the newly added sources either had international coverage or were based in Mexico with national coverage. This provides suggestive evidence that the appearance of new articles during this period did not occur as a result of adding sources that are likely to disproportionately cover particular regions, reducing concerns that

[61] I use weights proportional to the number of observations to account for imbalance in the panel data.

differential media coverage might have influenced the treatment effects.[62] As a result, the risk that media coverage biases my estimates during this crucial period is minimal, though the lack of sufficient pre-2008 data still presents a challenge for fully assessing pre-treatment trends in recall.

The main limitation of the above analysis is that it occurs at the state level, whereas treatment occurs at the municipality level. Because the recall analysis relies on state-level benchmarks, it cannot fully rule out the possibility of municipality-level shifts in coverage that coincide with treatment. I therefore complement the MCO-based checks with a municipality-level falsification test using non-organized crime related newspaper articles that mention a municipality in a given year as a dependent variable in my model. I use this as a proxy for general newspaper coverage in a region, as unlike the coverage that is used to detect CEs, it should not be affected by kingpin captures and deaths. If increased media coverage were driving the observed treatment effects, we would expect non-organized crime coverage to rise disproportionately in treated municipalities following kingpin captures and deaths. The results, shown in Appendix Figure A31 and Appendix Table A11, show that the treatment effect is negative but neither statistically nor economically significant, with an estimated decline of approximately 4% relative to the pre-treatment mean. This suggests that the observed treatment effects in the main result are driven by actual changes in criminal presence, rather than changes in reporting volume across locations.

5.3.3 Correlation of Enforcement Efforts with Treatment

As discussed in Sections 2.3 and 3.3, one possible threat to identification stems from the possibility that kingpin captures are caused by unequal investments in security personnel in treated municipalities and thus treatment is positively selected on. Ideally, I would test for this by testing whether municipality-level investments in security trended differently from a counterfactual weighted control in the periods prior to treatment. While this information is not publicly available due to security concerns, I obtained data on the number of Federal Police and National Guard officers deployed at the state-year level from Mexico's National Transparency Platform.[63] In Appendix Figure A32, I show that I am able to construct synthetic control groups using never-treated states that have trends that are parallel to the trend for treated states, suggesting that investment in enforcement was widespread and uncorrelated with treatment timing conditional on controls.

[62] A list of news sources included in the database by year is published by ProQuest.

[63] Mexico's Public Transparency Platform allows freedom of information requests and their responses to be viewed by the public. A request for this data was denied by the government on the basis of security concerns.

Similarly, police corruption presents a potential threat to identification because it may influence both where and when high-value targeting occurs. In contexts where criminal enterprises have ties to local officials, enforcement may be delayed or avoided, and these patterns could shift over time as political incentives or the criminal market structure changes. If such time-varying corruption influences both the timing of treatment and outcomes like violence or market structure, then the estimates presented earlier may be contaminated by selection bias.

While there are no direct measures of government corruption, one common proxy is the discrepancy between the number of homicides recorded in civil registry data—collected by medical examiners and published by INEGI—and the number of homicides reported by state-level prosecutors following an investigation, which is compiled by the Executive Secretariat of the National Public Security System (SESNSP) (Institute for Economics and Peace, 2016). Differences between these figures alone are not necessarily indicative of corruption: They can arise from cases where medical examiners and prosecutors reach different conclusions about the cause or intent of death (Hope, 2021). However, unusually large discrepancies between the INEGI and SESNSP homicide figures may be indicative of political incentives to under-report violence or institutional weaknesses that prevent cases from being formally investigated as homicides.

Ideally, to assess whether the timing of HVT is systematically related to local corruption, I would construct a panel of the per capita difference between INEGI-reported and SESNSP-reported homicides at the municipality-quarter level and test if it was correlated with the timing of kingpin captures or deaths. Evidence of such a relationship would raise concerns about selection bias, as treatment would no longer be plausibly independent of local reporting practices or enforcement capacity. Unfortunately, SESNSP only began publishing homicide data disaggregated by municipality and victim count starting in 2015, which coincides with the final kingpin capture in my sample. As a result, I cannot construct a panel over the relevant pre-treatment period.

Although the corruption proxy is only available after treatment, it can still provide suggestive evidence about whether corruption influenced the timing of high-value targeting. If time-varying corruption shaped when and where enforcement actions occurred, we would expect municipalities with larger homicide gaps in later years to have been treated earlier in the sample. I test this by regressing the total municipal-level per capita homicide gap between 2015 and 2019 on an indicator for whether or not the municipality had a kingpin capture or death during the event window. The results are reported in Appendix Table A12. The coefficient is small and statistically insignificant. The relationship remains so after con-

trolling for baseline criminal presence, which is positively correlated with both treatment and the homicide gap.[64] While this test may miss selection into HVT based on pre-existing corruption that does not persist post-treatment, the absence of a systematic relationship between early treatment and later gaps still provides suggestive evidence against persistent selection bias.

5.3.4 Inclusion of All Kingpins

In Section 3.3, I described how, inspired by Romer and Romer (2010), I use detailed narratives on the captures of each kingpin from news media to exclude those that threaten identification from the main analysis. This includes cases in which failed captures or misreporting generated possible anticipatory effects and captures that may have occurred as a result of a targeted operation that occurred in response to civilian massacres by Los Zetas. To demonstrate that these criteria are not selectively chosen to influence my findings, I present event study estimates for the main results with these events included in Appendix Figure A33 and in Appendix Tables A14 and A15. The results are robust to the inclusion of these kingpin captures or deaths. There is one small difference, which is that when all kingpins are included, there is evidence of some splintering, as the number of new splinters of CEs that were present in municipality m in the year before treatment increases relative to the control group. This makes sense, as the effect for splintering, while statistically insignificant, was imprecise and not close to zero in magnitude. Additionally, the results now include the death of Nazario Moreno González, whose organization, La Familia Michoacana, had reportedly splintered following his death (InSight Crime, 2014).[65] Still, the estimates for the new entry effect are larger in magnitude, so while I cannot rule out that splintering occurs, the entry of existing criminal enterprises is larger and more robust.

5.4 Imputation of Criminal Enterprises

As discussed previously, I assume that if there are 3 years or less of missing data two observations where a specific criminal enterprises is identified as operating in a specific location, then they operate in that location in the years in between as well. This is because it is likely that short lapses in coverage may be indicative of CEs not doing anything newsworthy, rather than exiting the territory. To ensure that this imputation is not driving the results, I com-

[64] In both models, the predicted coefficients are much smaller than the standard deviation of the per capita homicide gap, which is 272.13.

[65] I use the reported death of Moreno González in March of 2010 as the date of capture, as although he was falsely reported dead by the government in 2010, it was widely believed and Moreno González and according to reporting, had caused the organization to splinter.

pute my main results using the non-imputed data, which are reported in Appendix Figure A34 and Appendix Table A16. The non-imputed results are similar to the main results, suggesting that imputation is not driving the effects.

5.5 Alternative Treatment Geographies

In the main results, the unit of analysis is a Mexican municipality, with treated municipalities defined as those in which a kingpin capture or death occurred. Affected areas, however, may extend beyond these boundaries if spillovers occur.[66] To assess whether effects are more spatially diffuse, I re-assign treatment to include both the capture municipality and its immediate neighbors, aggregating outcomes and controls for these units into a single treated municipality.[67] Appendix Figure A35 and Appendix Tables A17-A18 present the results.

I continue to find a positive and statistically significant increase in the number of non-new splinter criminal enterprises not previously present and in the total number of criminal enterprises. The difference in estimated effects across specifications is modest and supports the interpretation that most entry occurs in the municipality where the capture took place. Consistent with this, the homicide rate effects lose significance in the spillover specification, likely because population increases outpace any rise in homicides.[68] Together, these results support the interpretation that the effects of high value targeting are highly localized and that the main specification, which excludes neighboring municipalities from both treated units and potential controls, appropriately captures the core area of market disruption.

6 Conclusion

This paper provides critical insights into the unintended consequences of targeting high-profile criminal leaders in Mexico. Leveraging a comprehensive longitudinal dataset derived from over 21 million full-text newspaper articles, I demonstrate that these interventions lead to significant market restructuring, marked by a 54% increase in the number of locally operating criminal enterprises in treated municipalities. This effect is largely driven by the entry of new competitors. Homicides rise by 32%, and while this violence is concentrated on

[66] In the main specification, neighboring municipalities and their neighbors are excluded from the set of possible controls as they are potentially affected by the policy.

[67] While it would be useful to define treatment based on where the affected criminal enterprise operated pre-capture, doing so would result in most municipalities being treated multiple times, which is incompatible with synthetic difference-in-differences.

[68] I do not standardize the number of criminal enterprises by population, since group presence reflects strategic territorial control rather than per capita exposure. A single group may operate in both densely and sparsely populated areas, and the potential associated harms (e.g., extortion, corruption, territorial conflict) are not experienced in proportion to population size.

organized crime members, it has significant implications for civilians. Although the entry effects are strongest in the period of the capture, there is a long-term increase in criminal presence and an escalation in violence for the two years following the capture or death of a kingpin.

These results suggest that capturing criminal leaders may temporarily disrupt criminal operations, but the strategy ultimately fails to produce lasting improvements in public safety. Instead of stabilizing the situation, the removal of a kingpin creates power vacuums that invite the entry of new criminal organizations. This influx of competitors is consistent with economic theory, where weakening a dominant player in the market lowers barriers to entry and incentivizes new firms to enter. In this case, the removal of a kingpin reduces the organizational strength of the incumbent, encouraging other groups to expand their influence. As new entrants vie for control, violence escalates, undermining public safety both in the short and long term.

The welfare implications are concerning. Communities where these captures take place bear the brunt of increased violence, facing both immediate harm and longer-term disruptions to social and economic stability. The rise in homicides further erodes public trust in law enforcement and the state. While capturing criminal leaders may achieve short-term tactical goals, the strategy ultimately exacerbates violence and undermines its broader objective of improving public safety.

Future research should explore whether similar dynamics occur in other contexts, particularly in counter-terrorism. Many countries have adopted leadership-targeting strategies against terrorist groups, yet little is known about whether these interventions produce comparable fragmentation and competition within these networks. For example, the power shifts following the U.S. withdrawal from Afghanistan, which saw the resurgence of groups like the Taliban and Al-Qaeda, offer a relevant parallel. Investigating whether leadership-targeting strategies have similar destabilizing effects in different settings could provide valuable insights into their broader implications.

A key contribution of this paper is the novel dataset I developed, which tracks the presence of criminal enterprises across Mexican municipalities. By using natural language processing techniques to extract detailed information newspaper articles, I capture a more granular view of criminal activity over time and across regions. This dataset enables an unprecedented examination of market structure and violence, offering a clearer understanding of how criminal organizations respond to law enforcement interventions. Beyond Mexico, this approach could be applied to study organized crime, insurgencies, or terrorist networks elsewhere. Additionally, the methodology could extend to other aspects of criminal enterprises, such as their

interaction with legal markets or local governance, providing new ways to understand how external shocks, such as leadership removals or state interventions, shape organized groups' behaviors and strategies.

References

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller, "Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program," *Journal of the American statistical Association*, 2010, *105* (490), 493–505.

Abbas, Mourad, Kamel Smaili, and Daoud Berkani, "Tr-classifier and knn evaluation for topic identification tasks," *International Journal on Information and Communication Technologies*, 2010, *3* (3), 10.

Alcocer, Marco, "Drug Wars, Organized Crime Expansion, and State Capture." PhD dissertation, UC San Diego 2023.

Arkhangelsky, Dmitry, Susan Athey, David A Hirshberg, Guido W Imbens, and Stefan Wager, "Synthetic difference-in-differences," *American Economic Review*, 2021, *111* (12), 4088–4118.

Arteaga, Carolina and Victoria Barone, "A manufactured tragedy: The origins and deep ripples of the opioid epidemic," Technical Report, Working paper 2022.

Astorga, Luis, "The limits of anti-drug policy in México," *International Social Science Journal*, 2001, *53* (169), 427–434.

Avram, Rachel, Eric J Koepcke, Alaa Moussawi, and Melissa Nuñez, "Do Cure Violence Programs Reduce Gun Violence? Evidence from New York City," *arXiv preprint arXiv:2406.02459*, 2024.

Battiston, Giacomo, Gianmarco Daniele, Marco Le Moglie, and Paolo Pinotti, "Fueling organized crime: The mexican war on drugs and oil thefts," 2022.

Brown, Ryan, Verónica Montalva, Duncan Thomas, and Andrea Velásquez, "Impact of violent crime on risk aversion: Evidence from the Mexican drug war," *Review of Economics and Statistics*, 2019, *101* (5), 892–904.

Calderón, Gabriela, Gustavo Robles, Alberto Díaz-Cayeros, and Beatriz Magaloni, "The beheading of criminal organizations and the dynamics of violence in Mexico," *Journal of Conflict Resolution*, 2015, *59* (8), 1455–1485.

Castillo, Juan Camilo, Daniel Mejía, and Pascual Restrepo, "Scarcity without leviathan: The violent effects of cocaine supply shortages in the mexican drug war," *Review of Economics and Statistics*, 2020, *102* (2), 269–286.

Cañete, José, Gabriel Chaperon, Rodrigo Fuentes, Jou-Hui Ho, Hojin Kang, and Jorge Pérez, "Spanish pre-trained BERT model and evaluation data," *PML4DC at ICLR 2020*, 2020.

Cedillo, Adela, *Operation Condor, the War on Drugs, and Counterinsurgency in the Golden Triangle (1977-1983)*, Helen Kellogg Institute for International Studies, 2021.

Ciccarone, Daniel, "The rise of illicit fentanyls, stimulants and the fourth wave of the opioid overdose crisis," *Current opinion in psychiatry*, 2021, *34* (4), 344–350.

Ciccia, Diego, "A Short Note on Event-Study Synthetic Difference-in-Differences Estimators," *arXiv preprint arXiv:2407.09565*, 2024.

Clarke, Damian, Daniel Pailanir, Susan Athey, and Guido Imbens, "Synthetic difference in differences estimation," *arXiv preprint arXiv:2301.11859*, 2023.

Coscia, Michele and Viridiana Rios, "Knowing where and how criminal organizations operate using web content," in "Proceedings of the 21st ACM international conference on Information and knowledge management" 2012, pp. 1412–1421.

CPJ, "Silence or death in Mexico's press: Crime, violence, and corruption are destroying the country's journalism," 2010.

Cruz, Ivan Lopez and Gustavo Torrens, "Hidden drivers of violence diffusion: Evidence from illegal oil siphoning in Mexico," *Journal of Economic Behavior & Organization*, 2023, *206*, 26–70.

Daniele, Gianmarco, Marco Le Moglie, and Federico Masera, "Pains, guns and moves: The effect of the US opioid epidemic on Mexican migration," *Journal of Development Economics*, 2023, *160*, 102983.

Dell, Melissa, "Trafficking networks and the Mexican drug war," *American Economic Review*, 2015, *105* (6), 1738–1779.

Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.

DHS, "US-mexico bi-national criminal proceeds study," *Department of Homeland Security*, 2010.

Fraioli, Paul, "Mexico's cartels and the rule of law," *Strategic Comments*, 2020, *26* (1), vii–ix.

Garcia, Klaifer, Pedro Shiguihara, and Lilian Berton, "Breaking news: Unveiling a new dataset for Portuguese news classification and comparative analysis of approaches," *Plos one*, 2024, *19* (1), e0296929.

Gavrilova, Evelina, Takuma Kamada, and Floris Zoutman, "Is legal pot crippling Mexican drug trafficking organisations? The effect of medical marijuana laws on US crime," *The Economic Journal*, 2019, *129* (617), 375–407.

Han, Xu, Tianyu Gao, Yankai Lin, Hao Peng, Yaoliang Yang, Chaojun Xiao, Zhiyuan Liu, Peng Li, Maosong Sun, and Jie Zhou, "More data, more relations, more context and more openness: A review and outlook for relation extraction," *arXiv preprint arXiv:2004.03186*, 2020.

Honnibal, Matthew, "Spacy's entity recognition model: Incremental parsing with bloom embeddings and residual CNNs," 2017.

Hope, Alejandro, "Para entender las cifras de homicidio," 2021. El Universal, Opinion Section.

Hurtado, Ignacio Rodriguez, "The Dose Response of Criminal Groups: Effects on Homicides and School Dropout," 2023.

InSight Crime, "Nazario Moreno González, 'El Chayo' 'El Más Loco'," *InSight Crime*, 2014. Accessed: 2025-05-20.

Institute for Economics and Peace, "Mexico Peace Index 2016: A Comprehensive Measure of Peace in Mexico," 2016. Accessed April 20, 2025.

Jiménez, Sergio Javier, "Retromará Calderón modelo colombiano contra narco," *El Universal*, Oct 2015.

Jones, Nathan, "Tijuana Cartel survives, despite decade-long onslaught," *InSight Crime*, Apr 2023.

Keefe, Patrick Radden, "Cocaine Incorporated," *The New York Times*, Jun 2012.

Lessing, Benjamin, "Making peace in drug wars: Crackdowns and cartels in Latin America," 2017.

Levitt, Steven D and Sudhir Alladi Venkatesh, "An economic analysis of a drug-selling gang's finances," *The quarterly journal of economics*, 2000, *115* (3), 755–789.

Lindo, Jason M and María Padilla-Romo, "Kingpin approaches to fighting crime and community violence: evidence from Mexico's drug war," *Journal of health economics*, 2018, *58*, 253–268.

Lopez, Itzel De Haro, "Avocados: Mexico's green gold. The US opioid crisis and its impact on Mexico's drug cartel violence," *Working Paper*, 2023.

Lupsha, Peter A, "Drug lords and narco-corruption: The players change but the game continues," 1991, pp. 177–195.

McDonnell, Patrick J, "How many people work for the Mexican Drug Cartels? researchers have an answer," *Los Angeles Times*, Sep 2023.

Medel, Monica and Francisco Thoumi, "Mexican drug "cartels"," *The Oxford handbook of organized crime*, 2014, pp. 196–218.

Mexico News Daily, "40 criminal groups behind drugs, extortion and murders in CDMX," *Mexico News Daily*, Jun 2020.

Miranda, Lester James, Ákos Kádár, Adriane Boyd, Sofie Van Landeghem, Anders Søgaard, and Matthew Honnibal, "Multi hash embeddings in SpaCy," *arXiv preprint arXiv:2212.09255*, 2022.

Molzhan, Cory, Octavio Rodriguez Ferreira, and David A. Shirk, "Drug violence in Mexico: Data and analysis through 2012," *Justice In Mexico*, Feb 2014.

National Archives and Records Administration, "Overview of Foreign Narcotics Kingpin Designation Act," *National Archives and Records Administration*, Apr 2009.

National Drug Intelligence Center, *National drug threat assessment*, National Drug Intelligence Center, 2000.

Neil, Daniel, Marwin Segler, Laura Guasch, Mohamed Ahmed, Dean Plumbley, Matthew Sellwood, and Nathan Brown, "Exploring deep recurrent models with reinforcement learning for molecule design," 2018.

Orozco-Aleman, Sandra and Heriberto Gonzalez-Lozano, "Drug violence and migration flows: Lessons from the Mexican drug war," *Journal of Human Resources*, 2018, *53* (3), 717–749.

Osorio, Javier, "Democratization and drug violence in Mexico," *Notre Dame: University of Notre Dame), unpublished typescript*, 2012.

_ , "The contagion of drug violence: spatiotemporal dynamics of the Mexican war on drugs," *Journal of Conflict Resolution*, 2015, *59* (8), 1403–1432.

_ **and Alejandro Beltran**, "Enhancing the detection of criminal organizations in Mexico using ML and NLP," in "2020 International Joint Conference on Neural Networks (IJCNN)" IEEE 2020, pp. 1–7.

Pardo, Bryce, Jirka Taylor, Jonathan P Caulkins, Beau Kilmer, Peter Reuter, and Bradley D Stein, "Future of fentanyl," 2019.

Phillips, Brian J, "How does leadership decapitation affect violence? The case of drug trafficking organizations in Mexico," *The Journal of Politics*, 2015, *77* (2), 324–336.

Prieto-Curiel, Rafael, Gian Maria Campedelli, and Alejandro Hope, "Reducing cartel recruitment is the only way to lower violence in Mexico," *Science*, 2023, *381* (6664), 1312–1316.

Ramos, Dulce, "Operación "Lince Norte", golpe contundente a los Zetas: Sedena," *Animal Politico*, Aug 2011.

Ramraj, S, R Arthi, Solai Murugan, and MS Julie, "Topic categorization of tamil news articles using pretrained word2vec embeddings with convolutional neural network," in "2020 International Conference on Computational Intelligence for Smart Power System and Sustainable Energy (CISPSSE)" IEEE 2020, pp. 1–4.

Ríos, Viridiana, "Why did Mexico become so violent? A self-reinforcing violent equilibrium caused by competition and enforcement," *Trends in organized crime*, 2013, *16*, 138–155.

Rizos, Georgios, Konstantin Hemker, and Björn Schuller, "Augment to prevent: short-text data augmentation in deep learning for hate-speech classification," in "Proceedings of the 28th ACM international conference on information and knowledge management" 2019, pp. 991–1000.

Romer, Christina D and David H Romer, "The macroeconomic effects of tax changes: estimates based on a new measure of fiscal shocks," *American economic review*, 2010, *100* (3), 763–801.

RSF, "2024 World Press Freedom index – journalism under political pressure," *Reporters Without Borders (RSF)*, 2024.

Saigal, Pooja and Vaibhav Khanna, "Multi-category news classification using Support Vector Machine based classifiers," *SN Applied Sciences*, 2020, *2* (3), 458.

Schippa, Camilla, "This is how much violence costs Mexico's economy," May 2016.

Seguridad Justicia y Paz, "Ranking 2021 de las 50 ciudades más violentas del mundo," *Seguridad, Justicia y Paz*, Mar 2022.

Signoret, Patrick, Marco Alcocer, Cecilia Farfan-Mendez, and Fernanda Sobrino, "Mapping Criminal Organizations in Mexico: State Panel 2007-2015," 2021.

Sobrino, Fernanda, "Mexican cartel wars: Fighting for the US opioid market," *Working Paper*, 2019.

SpaCy, "Spanish - spacy models documentation."

Stewart, Scott, "Tracking Mexico's cartels in 2019," Jan 2019.

Sviatschi, Maria Micaela, "Making a narco: Childhood exposure to illegal labor markets and criminal life paths," *Econometrica*, 2022, *90* (4), 1835–1878.

Trejo, Guillermo and Sandra Ley, "Why did drug cartels go to war in Mexico? Subnational party alternation, the breakdown of criminal protection, and the onset of large-scale violence," *Comparative Political Studies*, 2018, *51* (7), 900–937.

United Nations Office on Drugs and Crime, "Organized crime module 1 key issues: Defining organized crime," Apr 2018.

United States Department of Justice, "Drug Trafficking Organizations - National Drug Threat Assessment 2010 (unclassified)," Feb 2010.

UNODC, "2010 World Drug Report," *United Nations Office on Drugs and Crime*, 2010.

_ , "UNODC Homicide Statistics 2013," 2013. Accessed: 2024-10-29.

_ , "Global Study on Homicide 2019," 2019. Accessed: 2024-10-29.

U.S. Department of Justice, "Herfindahl-Hirschman Index," 2023.

U.S. Government Accountability Office, "Counternarcotics: Treasury reports some results from designating drug kingpins, but should improve information on agencies' expenditures," *Counternarcotics: Treasury Reports Some Results from Designating Drug Kingpins, but Should Improve Information on Agencies' Expenditures — U.S. GAO*, Jan 2019.

Utar, Hâle, "Firms and labor in times of violence: Evidence from the Mexican drug war," *Working Paper*, 2018.

Valle-Jones, Diego, "Strengths and weaknesses of crime data in Mexico," *Diego Valle-Jones's Blog*, Feb 2011.

Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, 2017, *30.*

Vela, David Saul, "Estos son los 40 grupos criminales que se disputan La Cdmx," Jun 2020.

Zúñiga, Vivian Mateos and David A Shirk, "The Impact of State-level U.S. Legalization Initiatives on Illegal Drug Flows," Jan 2022.

Tables and Figures

Figure 1. Monthly Homicide Rate in Mexico

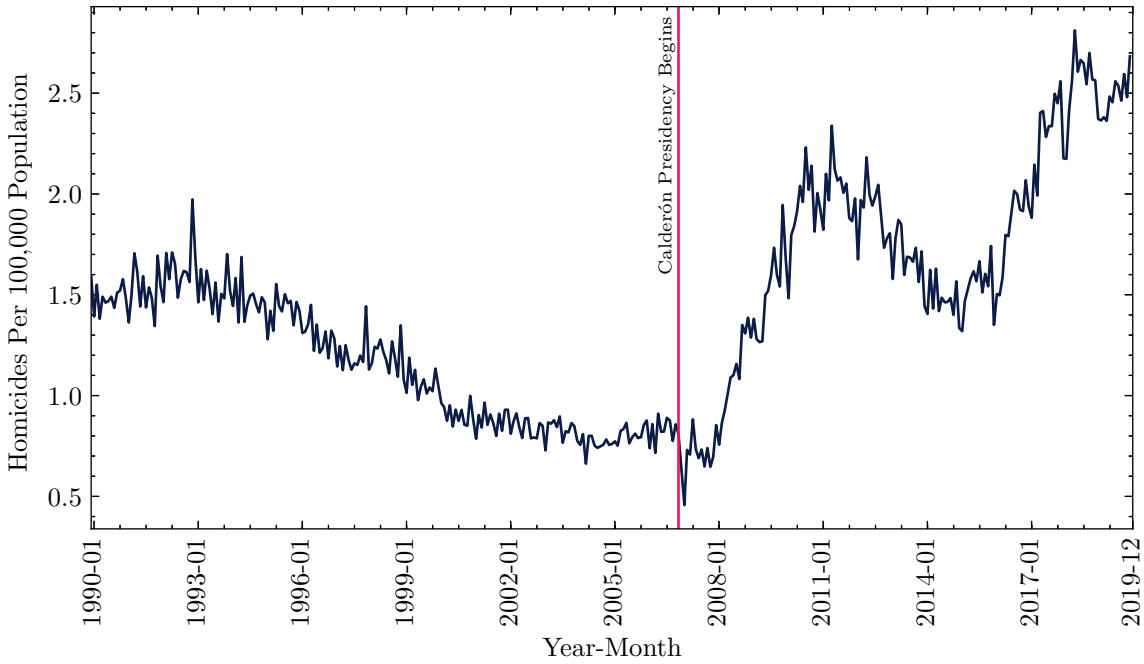

Notes: This figure plots the number of homicides per capita on a monthly basis in Mexico between January 1990 and December 2019. Homicide data is derived from death certificates, as reported by Mexico's statistical agency, the National Institute of Statistics and Geography (INEGI). It is converted to per-capita terms using census population data, which is linearly imputed between census years.

Figure 2. Number Of Kingpin Captures or Deaths By Year

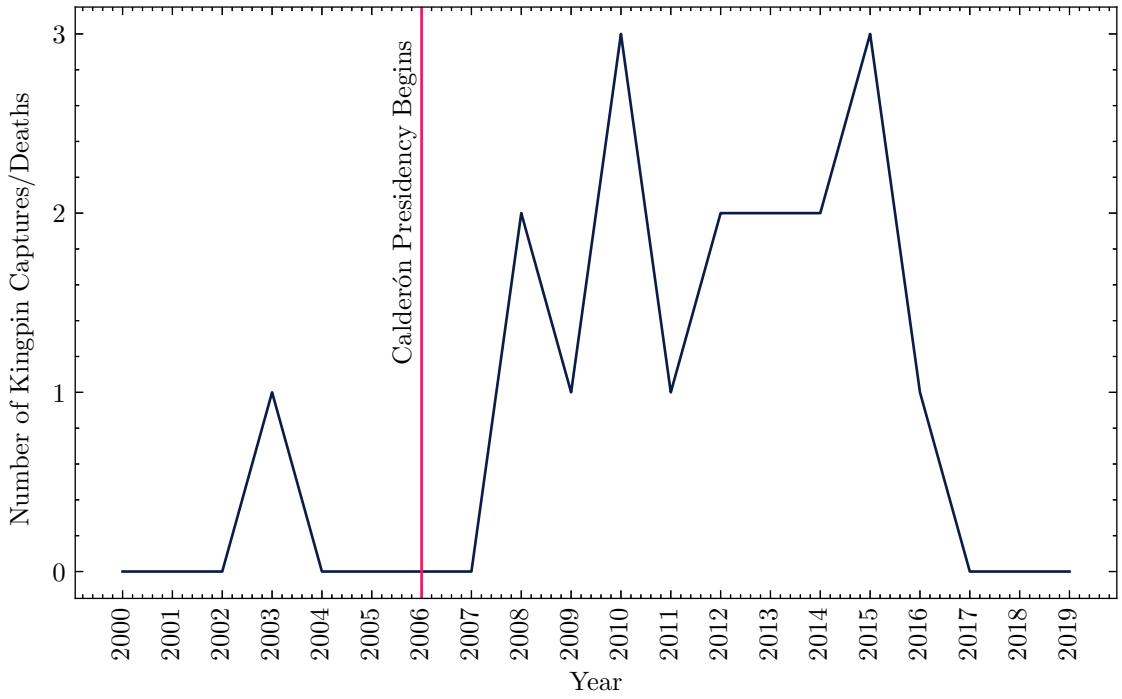

Notes: This figure plots the annual number of Mexican kingpins captured or killed by Mexican law enforcement. Kingpins are identified using the United States Office of Foreign Asset Control's Kingpin Sanctions List. Leaders are classified as kingpins if the leader is described in the news media as a top leader, or one of the top leaders, of a Mexican criminal enterprise. The figure includes all kingpins captures and deaths, including those that are excluded from the event study due to potential endogeneity, as discussed in Section 3.3.

Figure 3. Number of Municipalities With Kingpin Captures or Deaths By State

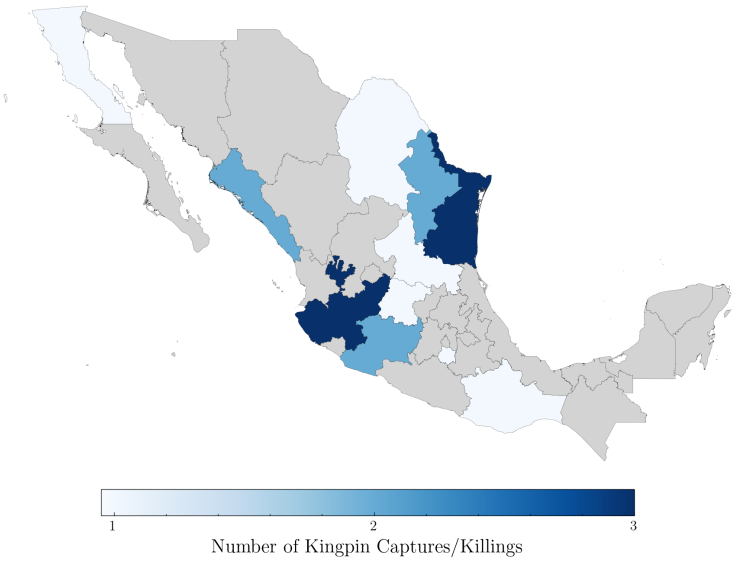

Notes: This figure plots the annual number of Mexican kingpins captured or killed by Mexican law enforcement at the state level between 2000 and 2019. Kingpins are identified using the United States Office of Foreign Asset Control's Kingpin Sanctions List. Leaders are classified as kingpins if the leader is described in the news media as a top leader, or one of the top leaders, of a Mexican criminal enterprise. The figure includes all kingpins captures and deaths, including those that are excluded from the event study due to potential endogeneity, as discussed in Section 3.3.

Figure 4. Average Annual Homicide Rate and Number of Operative Criminal Enterprises, By Municipality

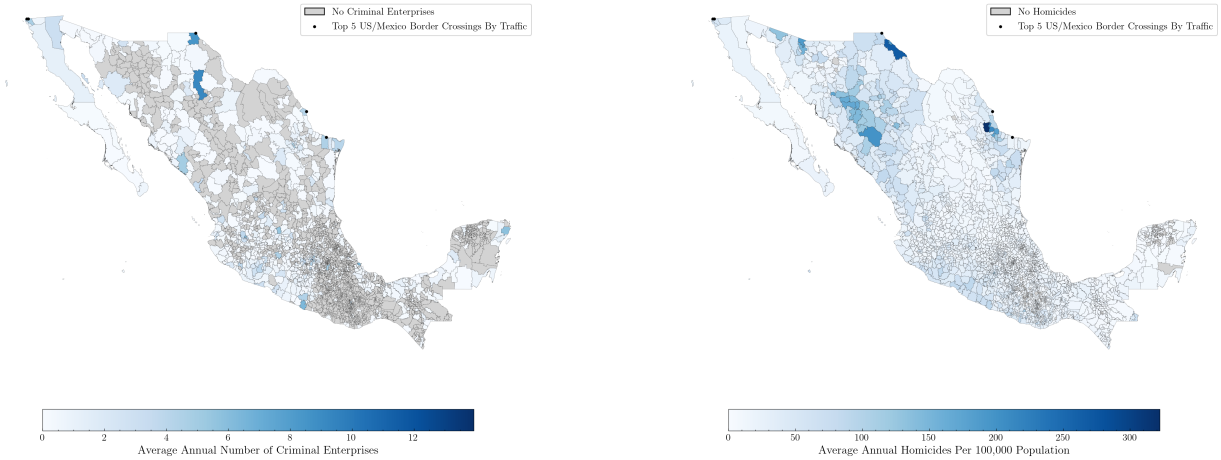

(a) Number of Criminal Enterprises

(b) Homicide Rate

Notes: Subfigure (a) plots the average annual number of criminal enterprises per municipality. The data is derived from newspaper articles in the ProQuest Latin American Newsstream using the data processing pipeline outlined in Section 3.1. Subfigure (a) plots the average annual homicide rate in Mexican municipalities between 2005 and 2019. Homicide data is derived from death certificates, as reported by Mexico's statistical agency, the National Institute of Statistics and Geography (INEGI). It is converted to per-capita terms using census population data, which is linearly imputed between census years. The figures also depict the locations of the 5 most frequently used United States border crossings along the U.S.-Mexico border (by either commercial or non-commercial vehicles) by any means of transport. They are identified using 2014 border traffic data from the United States Bureau of Transportation Statistics.

Figure 5. Average Annual Municipality-Level Market Concentration

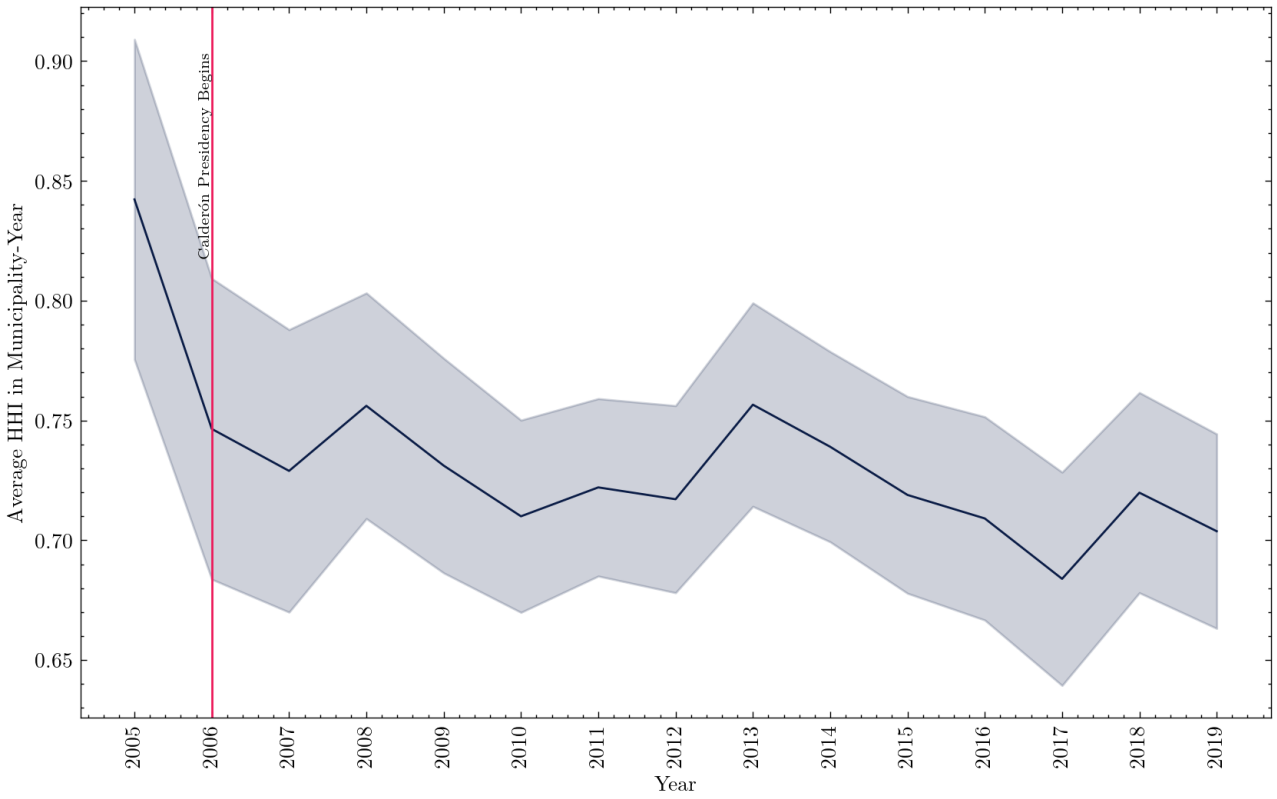

Notes: This figure plots the average annual Herfindahl-Hirschman Index (HHI) by municipality-year, representing market concentration across cartels and independent criminal organizations. The data is derived from newspaper articles in the ProQuest Latin American Newsstream using the data processing pipeline outlined in Section 3.1. The HHI is calculated as the sum of the squared market shares of criminal organizations in each municipality for a given year. I include cartels and independent criminal organizations (as defined in 3.1 in my HHI calculation, as factions may be working on behalf of a criminal organization and thus not competing with them. These shares are assumed to be evenly split among the organizations present in the municipality, meaning each group is assigned an equal share of the local market. The dataset is derived from newspaper articles processed through the methodology outlined in Section 3.1 and detailed in Appendix B. The vertical line at 2006 marks the beginning of the Calderón presidency, a period known for increased militarization of Mexico's drug policy. The shaded area around the average line reflects the 95% confidence interval based on the standard error of the mean.

Figure 6. Average Number of Criminal Enterprises Operating in Mexican Municipalities, by Year and Type

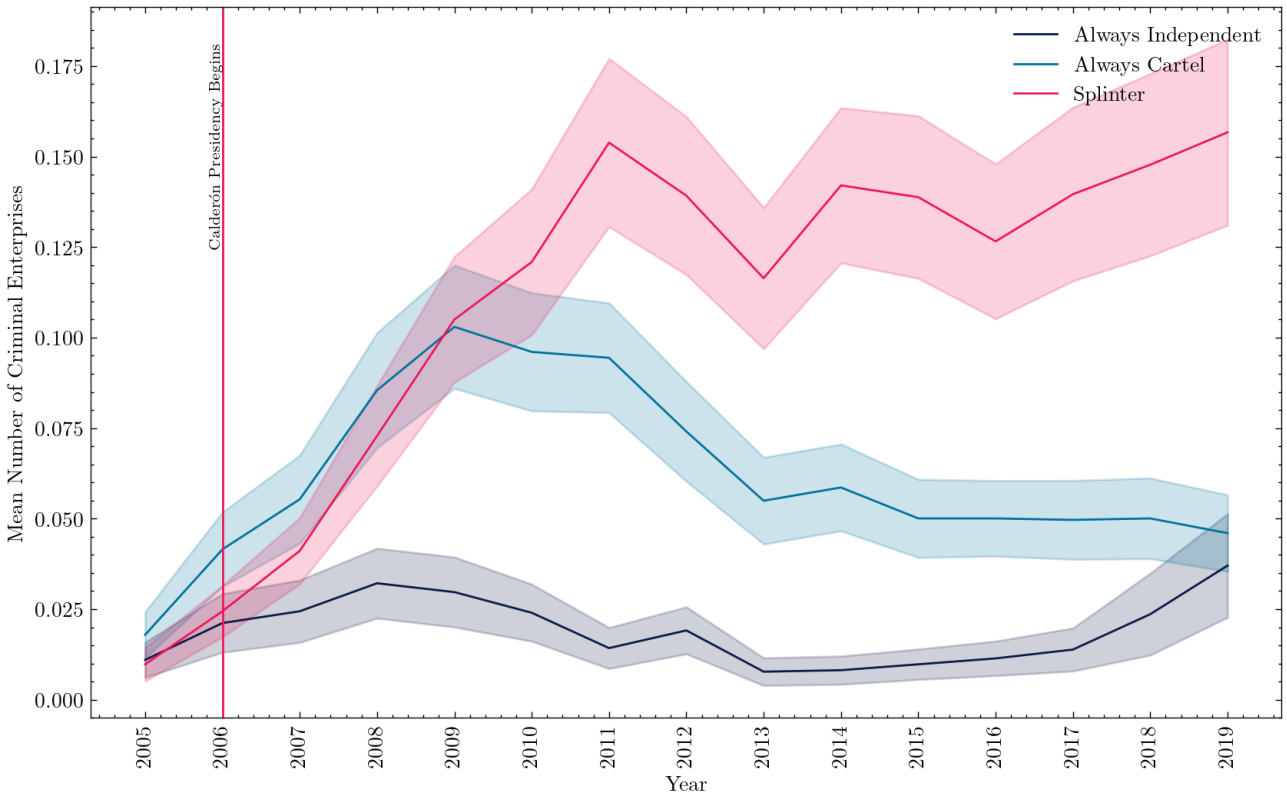

Notes: This figure plots the average number of unique splinter groups, as well as CEs that were always an independent group or always a cartel by year over the observation window (2005-2019). The data is derived from newspaper articles in the ProQuest Latin American Newsstream using the data processing pipeline outlined in Section 3.1. The shaded area around the average line reflects the 95% confidence interval based on the standard error of the mean.

Figure 7. Effect of Kingpin Captures and Deaths on Number of Criminal Enterprises

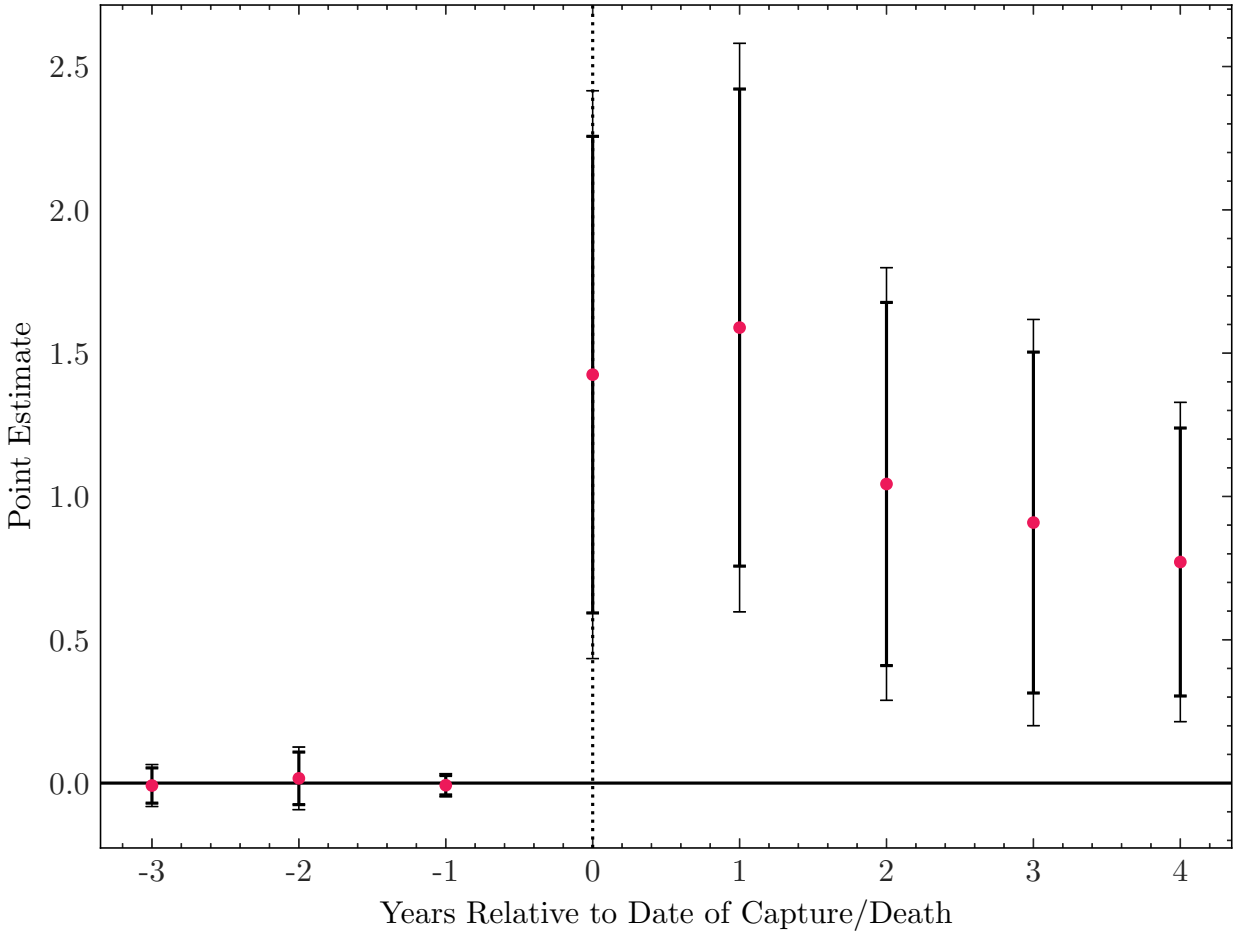

Notes: This figure plots estimates of the cohort-weighted average treatment effect on the treated $l \in \{-3, 4\}$ years relative to treatment, which is computed using the double difference event study estimator defined in Equation 10. The outcome is the number of criminal enterprises operating in municipality m in year t, which is residualized on the set of time-variant controls X^e_{mt}. The cohort specific ATT is computed as the difference between $\phi^e_{tr,l}$ and $\phi^e_{co,l}$. For each cohort e, $\phi^e_{tr,l}$ is the average difference between the residualized outcome in treated municipality m in period l, relative to the time-weighted pre-period mean for the residualized outcome. $\phi^e_{co,l}$ is the unit-weighted difference between the residualized outcome l years relative to treatment and the time-weighted pre-period mean for the never-treated control units. Cohort-level estimates are aggregated using weights proportional to the number of treated units in each cohort, as defined in Equation 7. The control set X^e_{mt}, discussed in Section 3.4, includes interactions between year indicator variables and each of the following: The suitability of municipality m for opium poppy cultivation, the suitability of municipality m for cannabis cultivation, and an indicator of whether municipality m contains a Pacific port with shipping container service. The unit and time weights are estimated following Arkhangelsky et al. (2021) as defined in Appendix Section C.1.The data is derived from newspaper articles in the ProQuest Latin American Newsstream using the data processing pipeline outlined in Section 3.1. Standard errors are obtained through block bootstrapping with 100 replications as outlined in Clarke et al. (2023) with sampling occurring at the state level to account for spatial correlation in treatment assignment. The bold vertical bands represent 90% confidence intervals. The other vertical bands represent 95% confidence intervals.

Figure 8. Effect of Kingpin Captures and Deaths on Homicides Per 100,000 Population

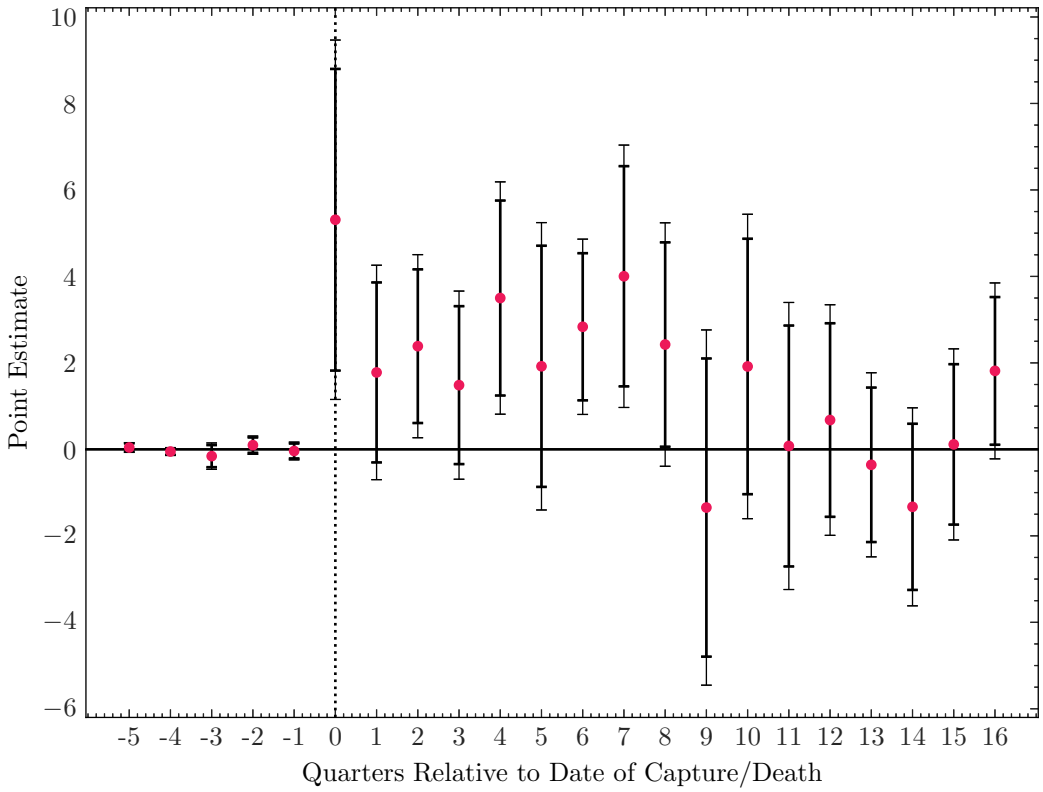

Notes: This figure plots estimates of the cohort-weighted average treatment effect on the treated $l \in \{-5, 16\}$ quarters relative to treatment, which is computed using the double difference event study estimator defined in Equation 10. The outcome is the per-capita quarterly homicide rate in municipality m in quarter t, which is residualized on the set of time-variant controls X^e_{mt}. Homicide data is derived from death certificates, as reported by Mexico's statistical agency, the National Institute of Statistics and Geography (INEGI). It is converted to per-capita terms using census population data, which is linearly imputed between census years. The cohort specific ATT is computed as the difference between $\phi^e_{tr,l}$ and $\phi^e_{co,l}$. For each cohort e, $\phi^e_{tr,l}$ is the average difference between the residualized outcome in treated municipality m in relative period l, relative to the time-weighted pre-period mean for the residualized outcome. $\phi^e_{co,l}$ is the unit-weighted difference between the residualized outcome l quarters relative to treatment and the time-weighted pre-period mean for the never-treated control units. Cohort-level estimates are aggregated using weights proportional to the number of treated units in each cohort, as defined in Equation 7. The control set X^e_{mt}, discussed in Section 3.4, includes interactions between a year-quarter indicator variable and each of the following: The suitability of municipality m for opium poppy cultivation, the suitability of municipality m for cannabis cultivation, and an indicator of whether municipality m contains a Pacific port with shipping container service. The unit and time weights are estimated following Arkhangelsky et al. (2021) as defined in Appendix Section C.1. Standard errors are obtained through block bootstrapping with 100 replications as outlined in Clarke et al. (2023) with sampling occurring at the state level to account for spatial correlation in treatment assignment. The bold vertical bands represent 90% confidence intervals. The other vertical bands represent 95% confidence intervals.

Table 1. Kingpin Capture/Death Dates

Kingpin Name	Organization	Municipality	Date	Event Type
Alfredo Beltrán Leyva	Cartel De Sinaloa	Culiacán, Sinaloa	January 21, 2008	Captured
Ramón Eduardo Arellano Félix	Cartel De Tijuana	Tijuana, Baja California	October 26, 2008	Captured
Arturo Beltrán Leyva	Beltran Leyva	Cuernavaca, Morelos	December 16, 2009	Killed
Ignacio Coronel Villareal	Cartel De Sinaloa	Zapopan, Jalisco	July 29, 2010	Killed
Flavio Méndez Santiago	Los Zetas	Villa de Etla, Oaxaca	January 18, 2011	Captured
Miguel Treviño Morales	Los Zetas	Anáhuac, Nuevo León	July 2013	Captured
Mario Ramirez Treviño	Cartel Del Golfo	Río Bravo, Tamaulipas	August 2013	Captured
Héctor Beltrán Leyva	Beltran Leyva	San Miguel de Allende, Guanajuato	October 1, 2014	Captured
Vicente Carillo-Fuentes	Cartel De Juarez	Torreón, Coahuila	October 9, 2014	Captured
Servando Gómez Martínez	Caballeros Templarios	Morelia, Michoacán	February 27, 2015	Killed
Abigael González Valencia	Cartel Jalisco Nueva Generacion	Puerto Vallarta, Jalisco	February 28, 2015	Captured
Omar Treviño Morales	Los Zetas	San Pedro Garza García, Nuevo León	March 4, 2015	Captured

Notes: This table reports details on kingpin captures/deaths that occurred during the event window (2008-2015) in chronological order. The leader's name, the organization to which they belong, the municipality in which they were captured, the date of capture, and whether they were killed or captured are listed for each leader. The table also lists the source(s) that report(s) on these details.

Table 2. Summary Statistics

	Treated Units	All of Mexico
Panel A: Municipality-Year Level		
Number of Criminal Enterprises	2.08	0.15
	(1.76)	(0.76)
Number of Cartels	1.83	0.10
	(1.58)	(0.52)
Number of Factions	0.08	0.02
	(0.28)	(0.21)
Number of Independents	0.17	0.02
	(0.45)	(0.20)
Population (in thousands)	504.44	40.29
	(474.69)	(122.77)
Number of Observations	36	32,361
Panel B: Municipality-Quarter Level		
Quarterly Homicides per 100,000 Population	4.89	2.64
	(5.35)	(12.11)
Number of Observations	60	107,814
Panel C: Municipality Characteristics		
Poppy Suitability	3.30	1.99
	(2.02)	(2.06)
Cannabis Suitability	78.60	18.64
	(260.56)	(98.52)
Number of Observations	12	2,163
Panel D: Captures		
Kingpin was Killed	0.25	
	(0.45)	
Number of Observations	12	
Number of Municipalities	12	2,163

Notes: This table reports summary statistics on the variables included in the SDID analysis. Column 1 shows the pre-treatment mean for treated units, while Column 2 shows the overall mean for all of Mexico (for all time periods). The table excludes municipalities covered by the sample restrictions, which are discussed in Section 3.5. Standard deviations appear in parentheses. The number of observations used for each calculation is listed at the bottom of each panel.

Table 3. Effect of Kingpin Captures and Deaths on Market Structure

	Pre-Treatment Mean		Level Change		% Change	
	(1)	(2)	(3)	(4)	(5)	(6)
Number of Criminal Enterprises	2.10	2.12	1.16***	1.15***	55.18***	54.26***
	(1.74)	(1.75)	(0.35)	(0.35)	(16.83)	(16.72)
Number of Non-New Splinter CEs Present in Year Before Treatment	0.88	0.89	0.08	0.08	9.14	9.15
	(1.16)	(1.18)	(0.11)	(0.11)	(12.60)	(12.39)
Number of Non-New Splinter CEs Not Present in Year Before Treatment	0.44	0.44	0.85***	0.84***	191.83***	192.12***
	(0.73)	(0.72)	(0.23)	(0.23)	(52.18)	(52.77)
Number of New Splinters of CEs Present in Year Before Treatment	0.39	0.40	0.24	0.24	60.77	59.10
	(0.67)	(0.67)	(0.17)	(0.17)	(43.82)	(42.88)
Number of New Splinters of CEs Not Present in Year Before Treatment	0.44	0.44	0.01	0.00	1.30	0.49
	(0.65)	(0.65)	(0.10)	(0.10)	(23.18)	(23.16)
Controls	No	Yes	No	Yes	No	Yes
Number of Cohorts	7	7	7	7	7	7
N Treated	12	12	12	12	12	12
N Control	2,151	2,151	2,151	2,151	2,151	2,151

Notes: The reported coefficients are estimated effects of kingpin captures or deaths on outcomes related to market structure, as estimated using Synthetic Difference-in-Differences and defined in Equation 6. The number of criminal enterprises includes CEs of all types (cartels, factions, and independent groups). I define new splinters as either new factions of a criminal enterprise (CE) or former factions that recently split from a CE (within the past 4 years). Columns (1) and (2) show the cohort-weighted and SDID weighted pre-treatment mean, with standard errors reported below. Column (1) uses control variables during the selection of the SDID weights, and column (2) does not. Columns (3) and (4) show the estimated treatment effects in levels with and without controls, respectively. The weighting process for the synthetic control units is described in Appendix C.1. Columns (5) and (6) show the percentage change in the outcome relative to the pre-treatment mean, weighted by the SDID weights, as defined in Equation 8. Controls include year indicators interacted with indicies measuring the suitability of municipality m for cultivating opium poppies and cannabis, respectively, as discussed in 3.4. Standard errors for coefficient estimates are obtained through block bootstrapping with 100 replications as outlined in Clarke et al. (2023). All regressions include year and municipality fixed effects. * p<0.1, ** p<0.05, *** p<0.01.

Table 4. Effect of Kingpin Captures and Deaths on Homicide Rates

	Pre-Treatment Mean		Level Change		% Change	
	(1)	(2)	(3)	(4)	(5)	(6)
Homicides Per 100,000 Population	4.89	5.04	2.46***	1.60**	50.24***	31.78**
	(5.28)	(5.40)	(0.70)	(0.70)	(14.39)	(13.96)
Homicides Per 100,000 Population (Males Aged 15-39)	2.55	3.07	2.02***	1.49***	79.09***	48.53***
	(2.90)	(3.29)	(0.50)	(0.50)	(19.65)	(16.33)
Homicides Per 100,000 Population (Excluding Males Aged 15-39)	2.22	2.00	1.16***	0.83***	52.06***	41.23***
	(3.81)	(3.73)	(0.29)	(0.29)	(13.05)	(14.47)
Controls	No	Yes	No	Yes	No	Yes
Number of Cohorts	9	9	9	9	9	9
N Treated	12	12	12	12	12	12
N Control	2,151	2,151	2,151	2,151	2,151	2,151

Notes: The reported coefficients are estimated effects of kingpin captures or deaths on outcomes related to homicide rates per 100,000 population, as estimated using Synthetic Difference-in-Differences and defined in Equation 6. I report estimates for the overall homicide rates, in addition to decomposing the results according to whether or not the victim is in a demographic group that is likely to be involved in organized crime. This includes males 15-39, as discussed by Calderón et al. (2015). Columns (1) and (2) show the cohort-weighted and SDID weighted pre-treatment mean, with standard errors reported below. Column (1) uses control variables during the selection of the SDID weights, and column (2) does not. Columns (3) and (4) show the estimated treatment effects in levels with and without controls, respectively. The weighting process for the synthetic control units is described in Appendix C.1. Columns (5) and (6) show the percentage change in the outcome relative to the pre-treatment mean, weighted by the SDID weights, as defined in Equation 8. Standard errors for coefficient estimates are obtained through block bootstrapping with 100 replications as outlined in Clarke et al. (2023). All regressions include year-quarter and municipality fixed effects. Additional controls include year-quarter indicators interacted with indicies measuring the suitability of municipality m for cultivating opium poppies and cannabis, respectively, as discussed in 3.4. Standard errors are clustered at the state level to account for correlation in treatment status within states. All regressions include year-quarter and municipality fixed effects. * p<0.1, ** p<0.05, *** p<0.01.

A Additional Tables and Figures

Appendix Figure A1. Monthly Homicides Per 100,000 Population in Municipalities with Capture or Death, by Kingpin

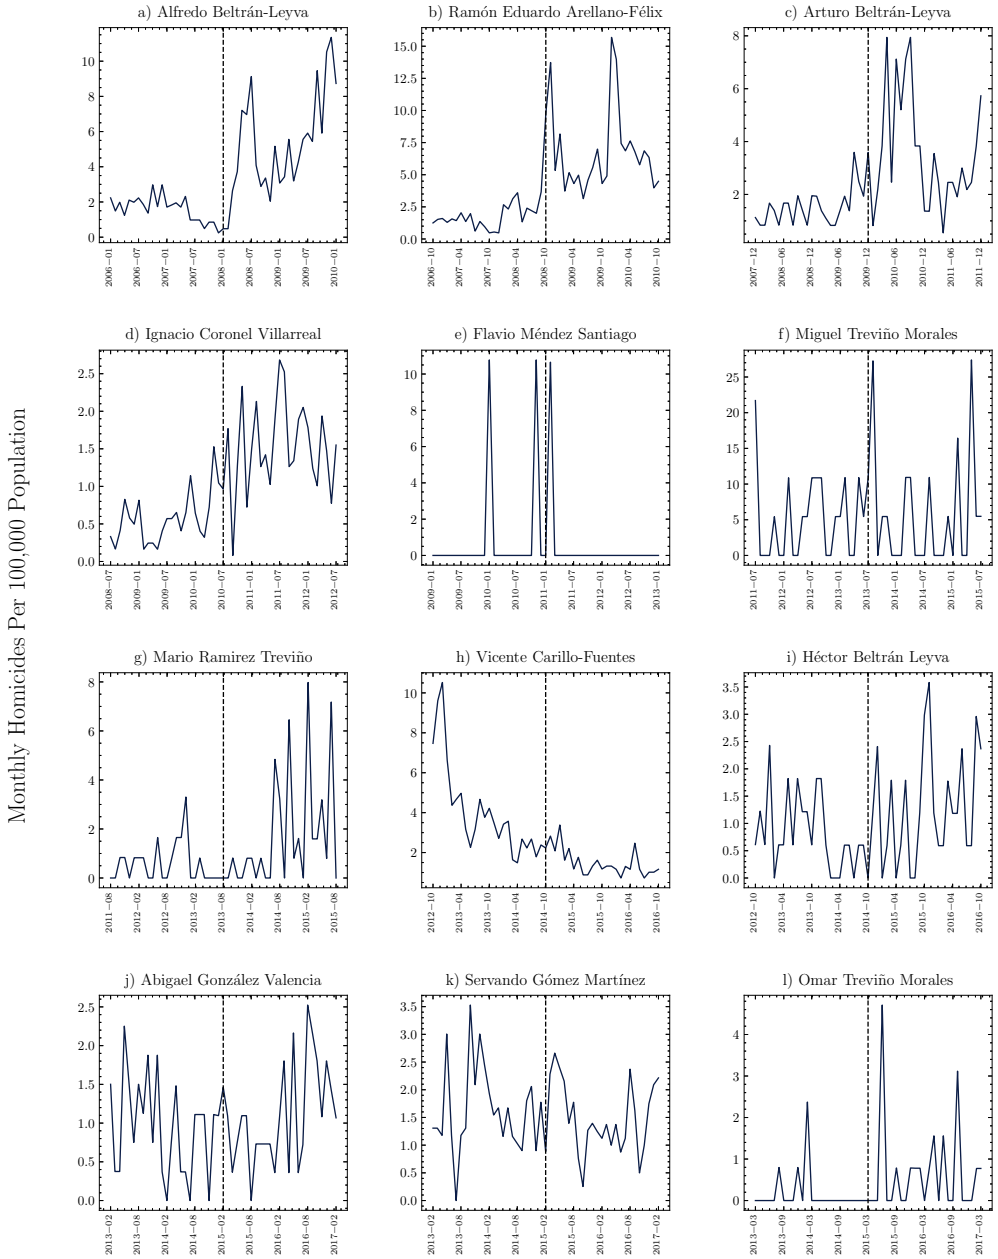

Notes: This figure plots the monthly homicide rate 24 months before and after each kingpin capture/death in the municipality of capture/death. The y-axis measures homicide rates per 100,000 population. The black dotted line is the date of capture/death. Homicide data is derived from death certificates, as reported by Mexico's statistical agency, the National Institute of Statistics and Geography (INEGI). The homicide rate is converted to per-capita terms using census population data, which is linearly imputed between census years.

Appendix Figure A2. Overall Homicide Rate Versus Drug-Related Killings

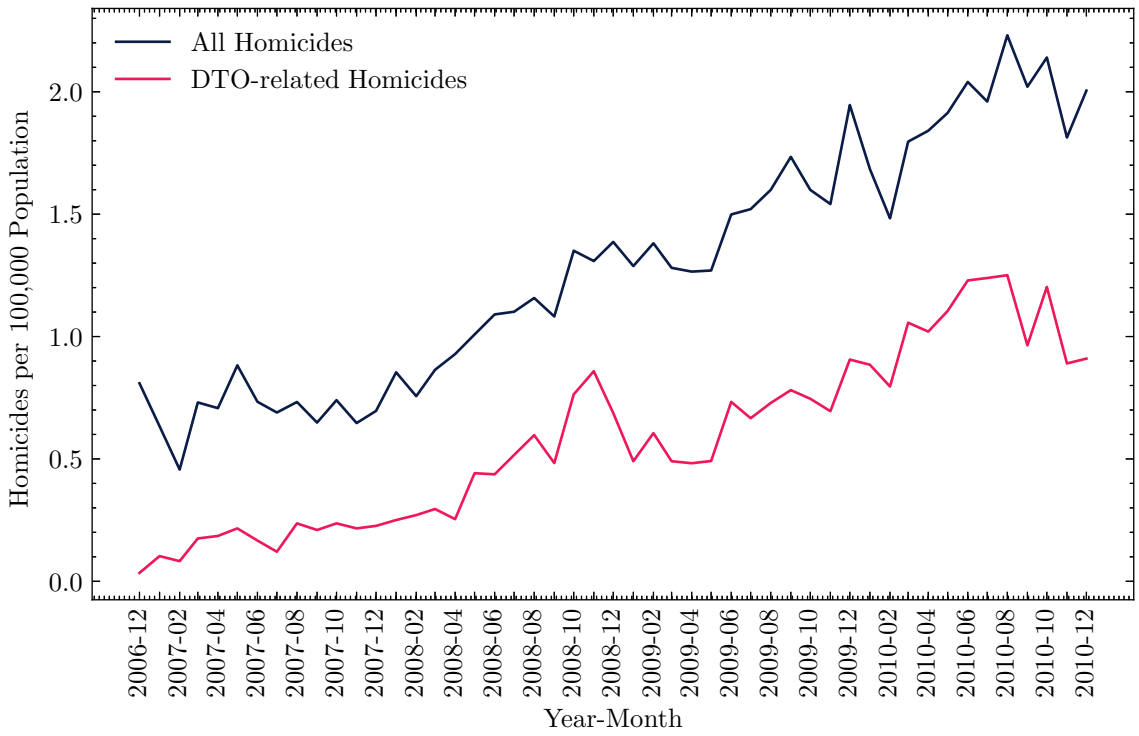

Notes: This figure compares the monthly drug-related homicide rate per 100,000 population against the number of overall homicides on a monthly basis from December 2006 to December 2010 (the period for which drug-related homicide data is available). Homicides are classified as drug-related if one civilian killed another civilian with at least one of those involved in the drug trade. Drug-related homicide data is from the National Council of Public Security. Overall homicide data is derived from death certificates, as reported by Mexico's statistical agency, the National Institute of Statistics and Geography (INEGI). The homicide rate is converted to per-capita terms using census population data, which is linearly imputed between census years.

Appendix Figure A3. Average Municipal Criminal Presence by Municipality and Time Period

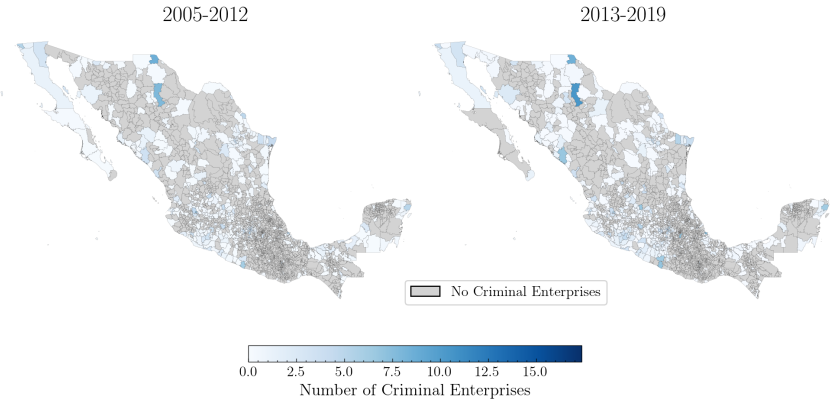

Notes: This figure shows the average number of criminal enterprises operating in each municipality during the periods 2005–2012 and 2014–2019. For each municipality, the average number of enterprises is calculated separately for each period. Darker shades indicate municipalities with a higher average number of criminal enterprises. The data is derived from newspaper articles in the ProQuest Latin American Newsstream using the data processing pipeline outlined in Section 3.1.

Appendix Figure A4. Adjusted Average Annual Municipality-Level Market Concentration

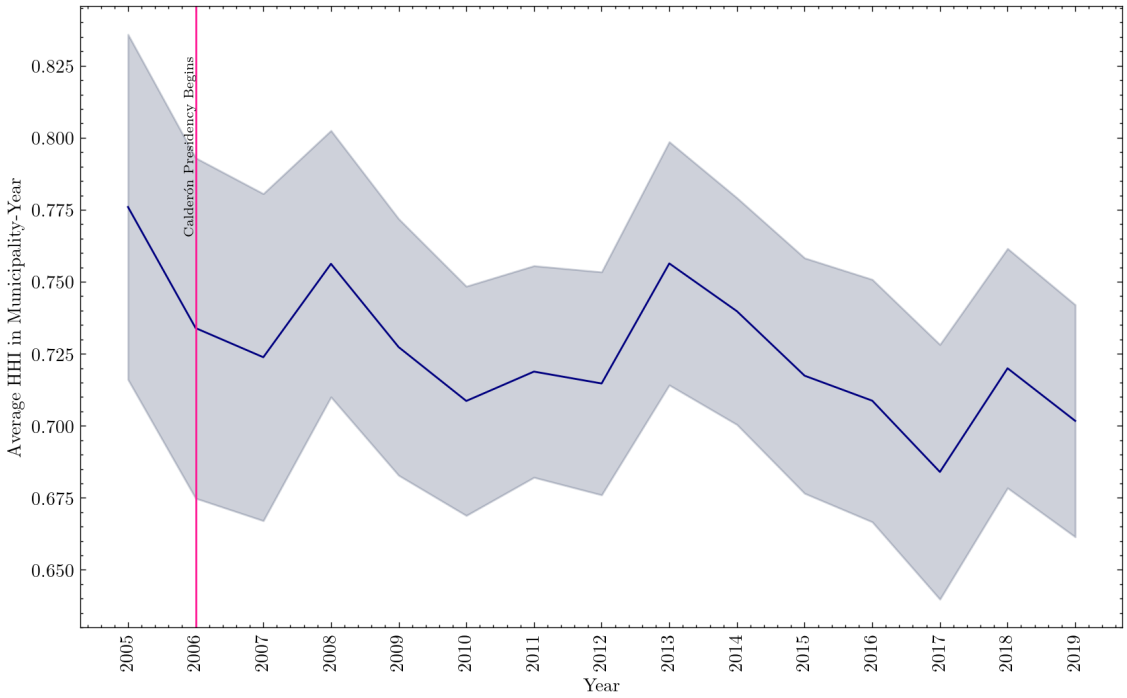

Notes: This figure plots the adjusted average annual Herfindahl-Hirschman Index (HHI) by municipality-year, representing market concentration across cartels and independent criminal organizations. This figure adjusts for the role that newspaper coverage plays in detecting CEs using the process described in Section 3.6. The data is derived from newspaper articles in the ProQuest Latin American Newsstream using the data processing pipeline outlined in Section 3.1. The HHI is calculated as the sum of the squared market shares of criminal organizations in each municipality for a given year. I include cartels, independent criminal organizations, and splinter groups (as defined in 3.1 in my HHI calculation, and exclude factions that are known to be loyal to a criminal organization. These shares are assumed to be evenly split among the organizations present in the municipality, meaning each group is assigned an equal share of the local market. Municipalities that do not have a criminal presence are excluded from the analysis, as these municipalities do not have a meaningful HHI. The dataset is derived from newspaper articles processed through the methodology outlined in Section 3.1 and detailed in Appendix B. The vertical line at 2006 marks the beginning of the Calderón presidency, a period known for increased militarization of Mexico's drug policy. The shaded area around the average line reflects the 95% confidence interval based on the standard error of the mean.

Appendix Figure A5. Adjusted Average Municipal Criminal Presence by Municipality and Time Period

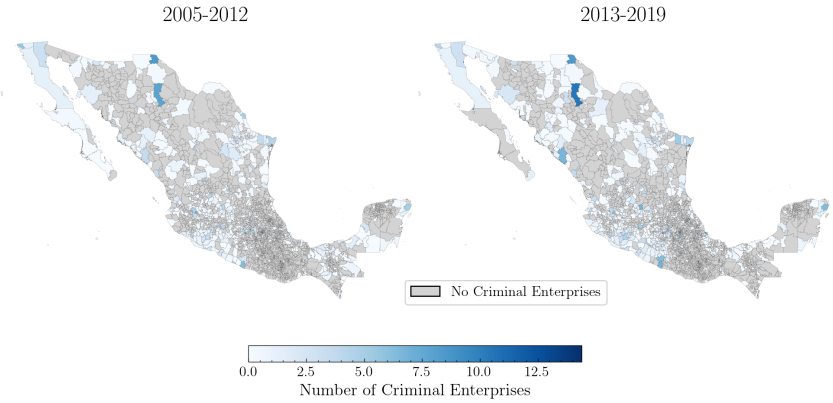

Notes: This figure shows the adjusted average number of criminal enterprises operating in each municipality during the periods 2005–2012 and 2014–2019. This figure adjusts for the role that newspaper coverage plays in detecting CEs using the process described in Section 3.6. For each municipality, the average number of enterprises is calculated separately for each period. Darker shades indicate municipalities with a higher average number of criminal enterprises. The data is derived from newspaper articles in the ProQuest Latin American Newsstream using the data processing pipeline outlined in Section 3.1.

Appendix Figure A6. Adjusted Average Number of Criminal Enterprises Operating in Mexican Municipalities with a Positive Criminal Presence, by Year and Type

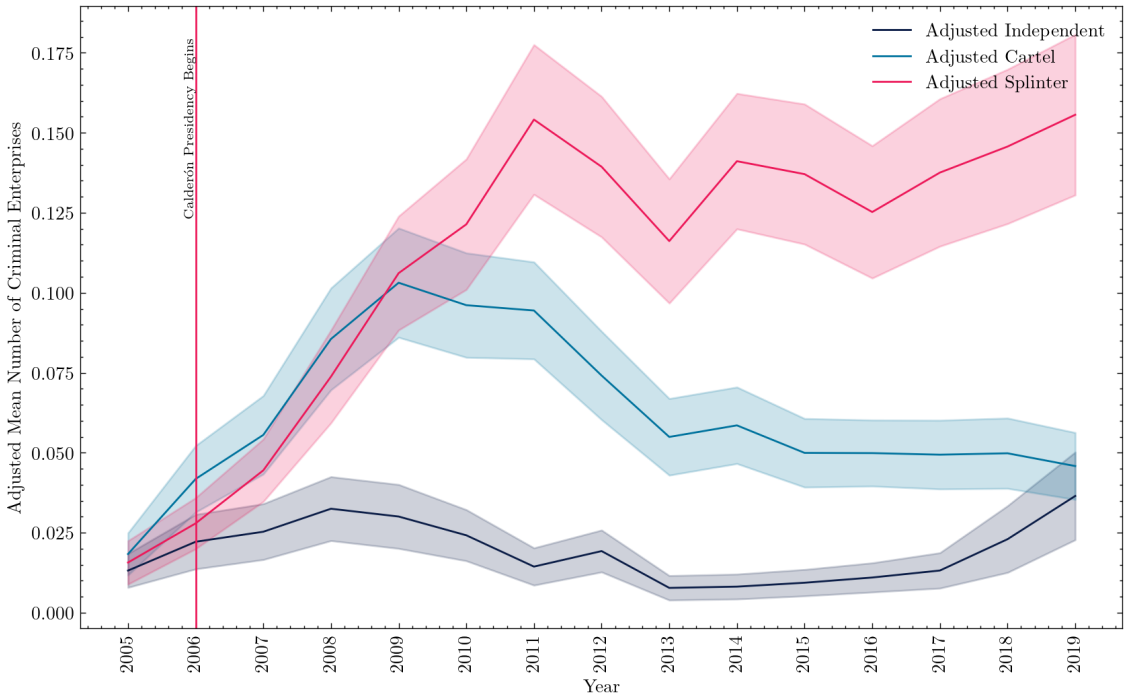

Notes: This figure plots the average number of unique splinter groups, as well as CEs that were always an independent group or always a cartel by year over the observation window (2005-2019). This figure adjusts for the role that newspaper coverage plays in detecting CEs using the process described in Section 3.6. The data is derived from newspaper articles in the ProQuest Latin American Newsstream using the data processing pipeline outlined in Section 3.1. The shaded area around the average line reflects the 95% confidence interval based on the standard error of the mean.

Appendix Figure A7. Effect of Kingpin Captures/Deaths on Number of Non-New Splinter Criminal Enterprises, According to CE Was Present In Year Prior to Capture/Death

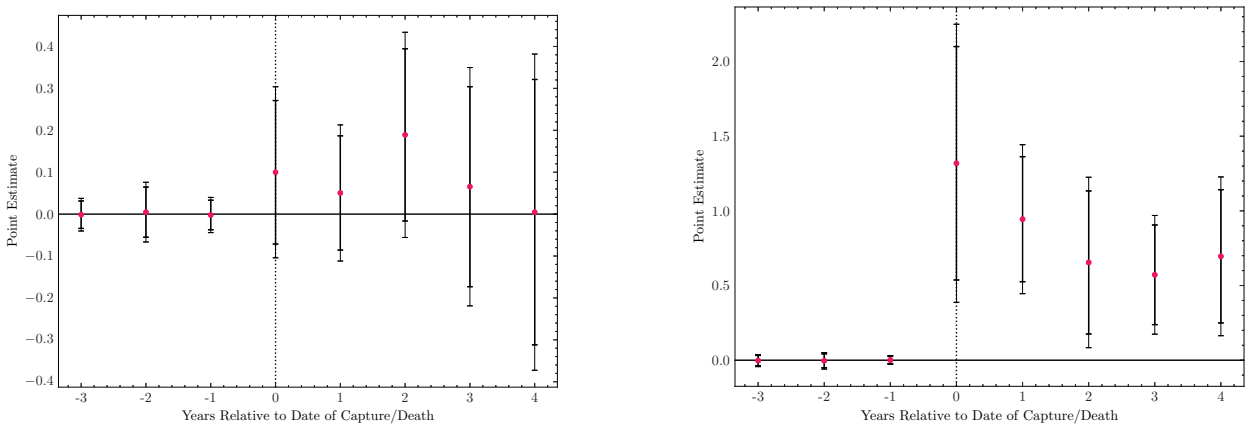

Present Not Present

Notes: This figure plots estimates of the cohort-weighted average treatment effect on the treated $l \in \{-3, 4\}$ years relative to treatment, which is computed using the double difference event study estimator defined in Equation 10. The outcomes in panels (a) and (b) are the number of non-new splinter criminal enterprises of a CE that was present and not present in the year before treatment in municipality m in year t, respectively, which are residualized on the set of time-variant controls X^e_{mt}. I define new splinters as either new factions of a criminal enterprise (CE) or former factions that recently split from a CE (within the past 4 years). The cohort specific ATT is computed as the difference between $\phi^e_{tr,l}$ and $\phi^e_{co,l}$. For each cohort e, $\phi^e_{tr,l}$ is the average difference between the residualized outcome in treated municipality m in period l, relative to the time-weighted pre-period mean for the residualized outcome. $\phi^e_{co,l}$ is the unit-weighted difference between the residualized outcome l years relative to treatment and the time-weighted pre-period mean for the never-treated control units. Cohort-level estimates are aggregated using weights proportional to the number of treated units in each cohort, as defined in Equation 7. The control set X^e_{mt}, discussed in Section 3.4, includes interactions between year indicator variables and each of the following: The suitability of municipality m for opium poppy cultivation, the suitability of municipality m for cannabis cultivation, and an indicator of whether municipality m contains a Pacific port with shipping container service. The unit and time weights are estimated following Arkhangelsky et al. (2021) as defined in Appendix Section C.1. The data is derived from newspaper articles in the ProQuest Latin American Newsstream using the data processing pipeline outlined in Section 3.1. Standard errors are obtained through block bootstrapping with 100 replications as outlined in Clarke et al. (2023) with sampling occurring at the state level to account for spatial correlation in treatment assignment. The bold vertical bands represent 90% confidence intervals. The other vertical bands represent 95% confidence intervals.

Appendix Figure A8. Effect of Kingpin Captures/Deaths on Number of New Splinter Criminal Enterprises
According to if Parent CE Was Present In Year Prior to Capture/Death

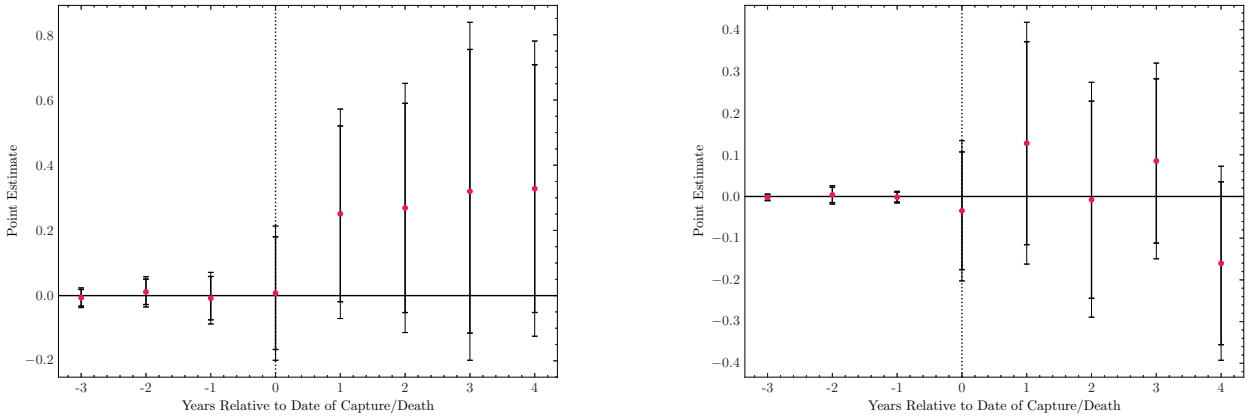

Present | Not Present

Notes: This figure plots estimates of the cohort-weighted average treatment effect on the treated $l \in \{-3, 4\}$ years relative to treatment, which is computed using the double difference event study estimator defined in Equation 10. The outcomes in panels (a) and (b) are the number of new splinters of CEs present and not present in the year before treatment in municipality m in year t, respectively, which are residualized on the set of time-variant controls X^e_{mt}. I define new splinters as either new factions of a criminal enterprise (CE) or former factions that recently split from a CE (within the past 4 years). The cohort specific ATT is computed as the difference between $\phi^e_{tr,l}$ and $\phi^e_{co,l}$. For each cohort e, $\phi^e_{tr,l}$ is the average difference between the residualized outcome in treated municipality m in period l, relative to the time-weighted pre-period mean for the residualized outcome. $\phi^e_{co,l}$ is the unit-weighted difference between the residualized outcome l years relative to treatment and the time-weighted pre-period mean for the never-treated control units. Cohort-level estimates are aggregated using weights proportional to the number of treated units in each cohort, as defined in Equation 7. The control set X^e_{mt}, discussed in Section 3.4, includes interactions between year indicator variables and each of the following: The suitability of municipality m for opium poppy cultivation, the suitability of municipality m for cannabis cultivation, and an indicator of whether municipality m contains a Pacific port with shipping container service. The unit and time weights are estimated following Arkhangelsky et al. (2021) as defined in Appendix Section C.1. The data is derived from newspaper articles in the ProQuest Latin American Newsstream using the data processing pipeline outlined in Section 3.1. Standard errors are obtained through block bootstrapping with 100 replications as outlined in Clarke et al. (2023) with sampling occurring at the state level to account for spatial correlation in treatment assignment. The bold vertical bands represent 90% confidence intervals. The other vertical bands represent 95% confidence intervals.

Appendix Figure A9. Effect of Kingpin Captures and Deaths on Homicides Per 100,000 Population, According to Victim's Likely Criminal Status

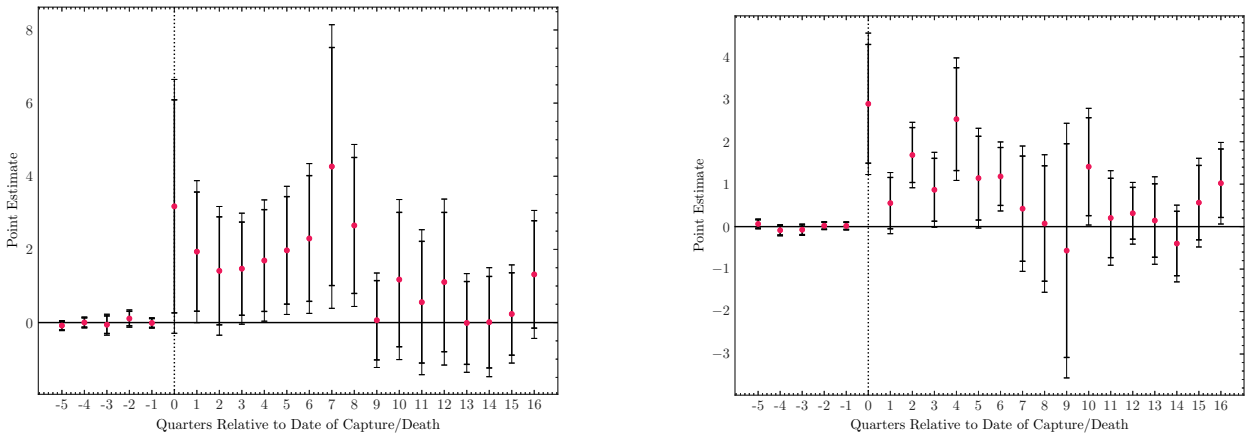

(a) Likely CE Member: Males Aged 15-39

(b) Unlikely CE Member: Remainder of Population

Notes: This figure plots estimates of the cohort-weighted average treatment effect on the treated $l \in \{-5, 16\}$ quarters relative to treatment, which is computed using the double difference event study estimator defined in Equation 10. The outcomes in panels (a) and (b) are the rate of homicides per 100,000 population in municipality m in quarter t for males aged 15-39 and the remainder of the population, respectively, which are residualized on the set of time-variant controls X^e_{mt}. The age categories are the same as those chosen by Calderón et al. (2015), who show that this variation most closely resembles data on drug-related homicides for years in which it was available. Homicide data is derived from death certificates, as reported by Mexico's statistical agency, the National Institute of Statistics and Geography (INEGI). It is converted to per-capita terms using census population data, which is linearly imputed between census years. The cohort specific ATT is computed as the difference between $\phi^e_{tr,l}$ and $\phi^e_{co,l}$. For each cohort e, $\phi^e_{tr,l}$ is the average difference between the residualized outcome in treated municipality m in period l, relative to the time-weighted pre-period mean for the residualized outcome. $\phi^e_{co,l}$ is the unit-weighted difference between the residualized outcome l quarters relative to treatment and the time-weighted pre-period mean for the never-treated control units. Cohort-level estimates are aggregated using weights proportional to the number of treated units in each cohort, as defined in Equation 7. The control set X^e_{mt}, discussed in Section 3.4, includes interactions between a year-quarter indicator variable and each of the following: The suitability of municipality m for opium poppy cultivation, the suitability of municipality m for cannabis cultivation, and an indicator of whether municipality m contains a Pacific port with shipping container service. The unit and time weights are estimated following Arkhangelsky et al. (2021) as defined in Appendix Section C.1. Standard errors are obtained through block bootstrapping with 100 replications as outlined in Clarke et al. (2023) with sampling occurring at the state level to account for spatial correlation in treatment assignment. The bold vertical bands represent 90% confidence intervals. The other vertical bands represent 95% confidence intervals.

Appendix Figure A10. Number of Criminal Enterprises Over Time for Municipalities with Kingpin Captures and Synthetic Control Unit, by Treatment Date

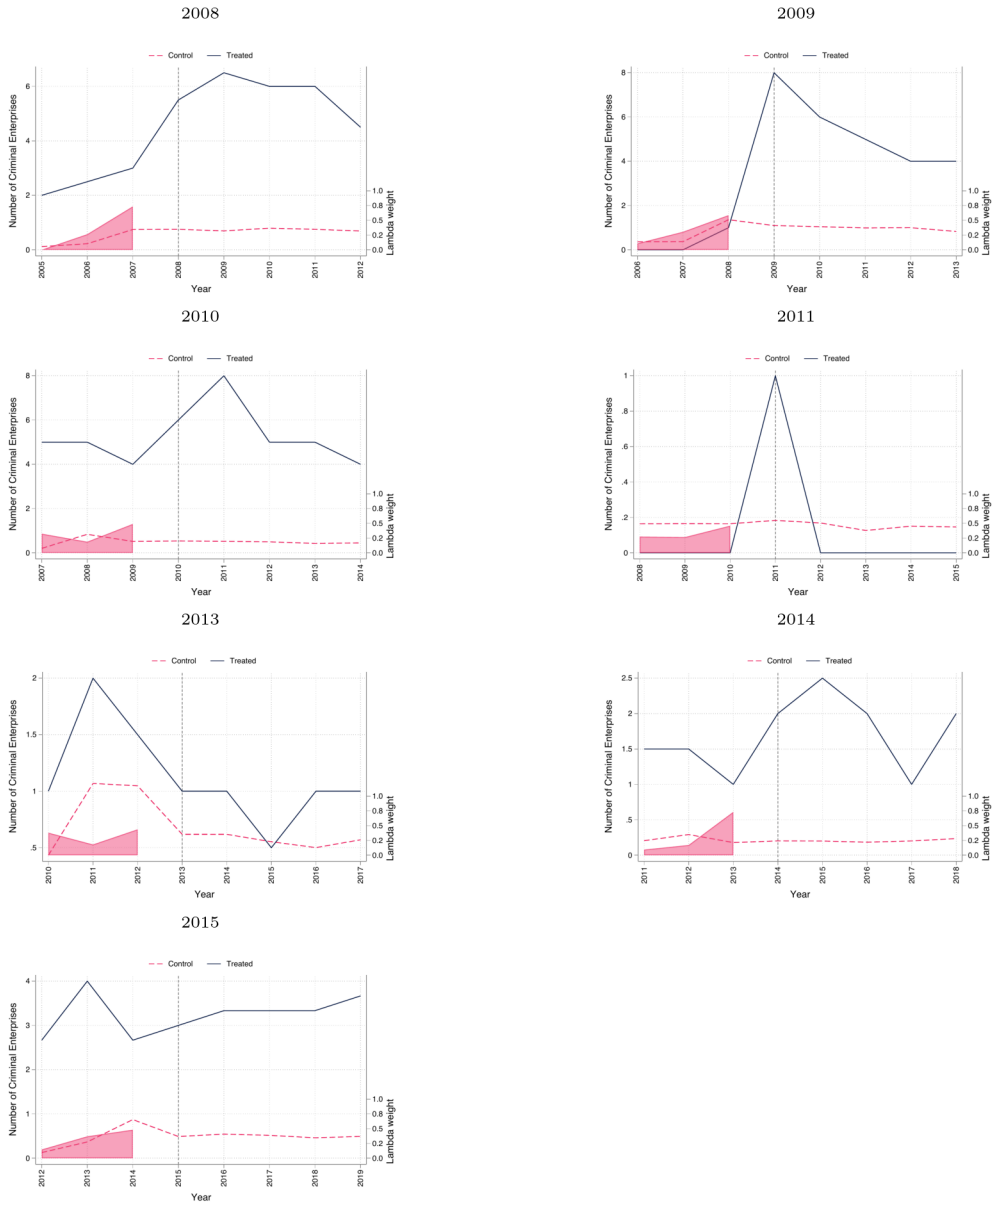

Notes: This figure plots, for each treatment date, the average number of criminal enterprises for municipalities with kingpin captures and deaths in year t, along with the analogous ω weighted average for the synthetic control unit. The λ weights used to average pretreatment time periods are at the bottom of the graph. The control set X^e_{mt}, discussed in Section 3.4, includes interactions between year indicator variables and each of the following: The suitability of municipality m for opium poppy cultivation, the suitability of municipality m for cannabis cultivation, and an indicator of whether municipality m contains a Pacific port with shipping container service.

Appendix Figure A11. Number of Non-New Splinter CEs That Were Present in Year Before Treatment Over Time for Municipalities with Kingpin Captures and Synthetic Control Unit, by Treatment Date

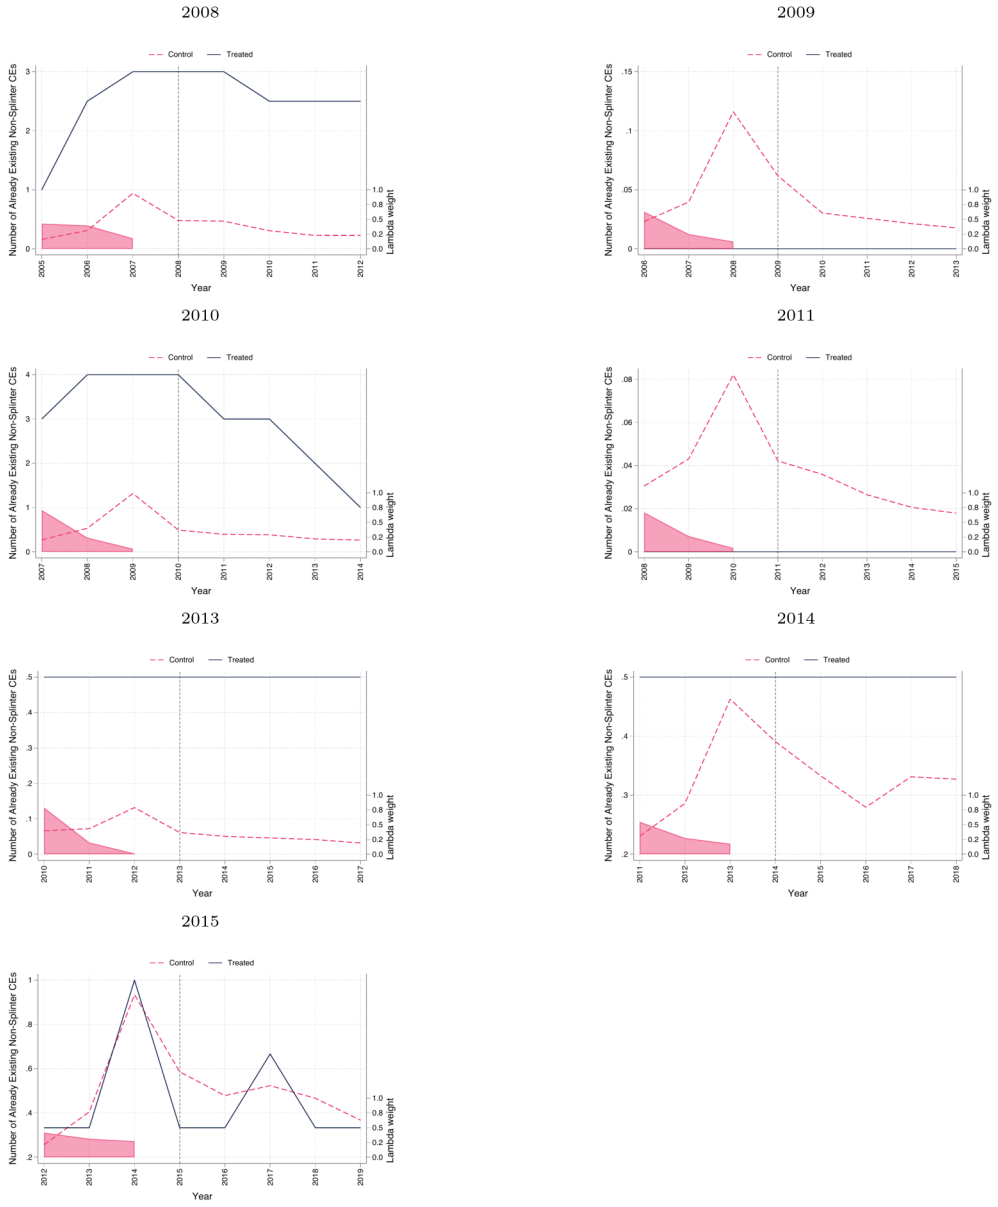

Notes: This figure plots, for each treatment date, the average number of non-new splinter criminal enterprises that were present in the year before treatment for municipalities with kingpin captures and deaths in year t, along with the analogous ω weighted average for the synthetic control unit. The λ weights used to average pretreatment time periods are at the bottom of the graph. The weight construction accounts for time variant controls, as discussed in Appendix Section C.1, however, the plots apply these weights to the non-residualized outcomes to make the graphs more interpretable. The control set X^e_{mt}, discussed in Section 3.4, includes interactions between year indicator variables and each of the following: The suitability of municipality m for opium poppy cultivation, the suitability of municipality m for cannabis cultivation, and an indicator of whether municipality m contains a Pacific port with shipping container service.

Appendix Figure A12. Number of Non-New Splinter CEs That Were Not Present in Year Before Treatment Over Time for Municipalities with Kingpin Captures and Synthetic Control Unit, by Treatment Date

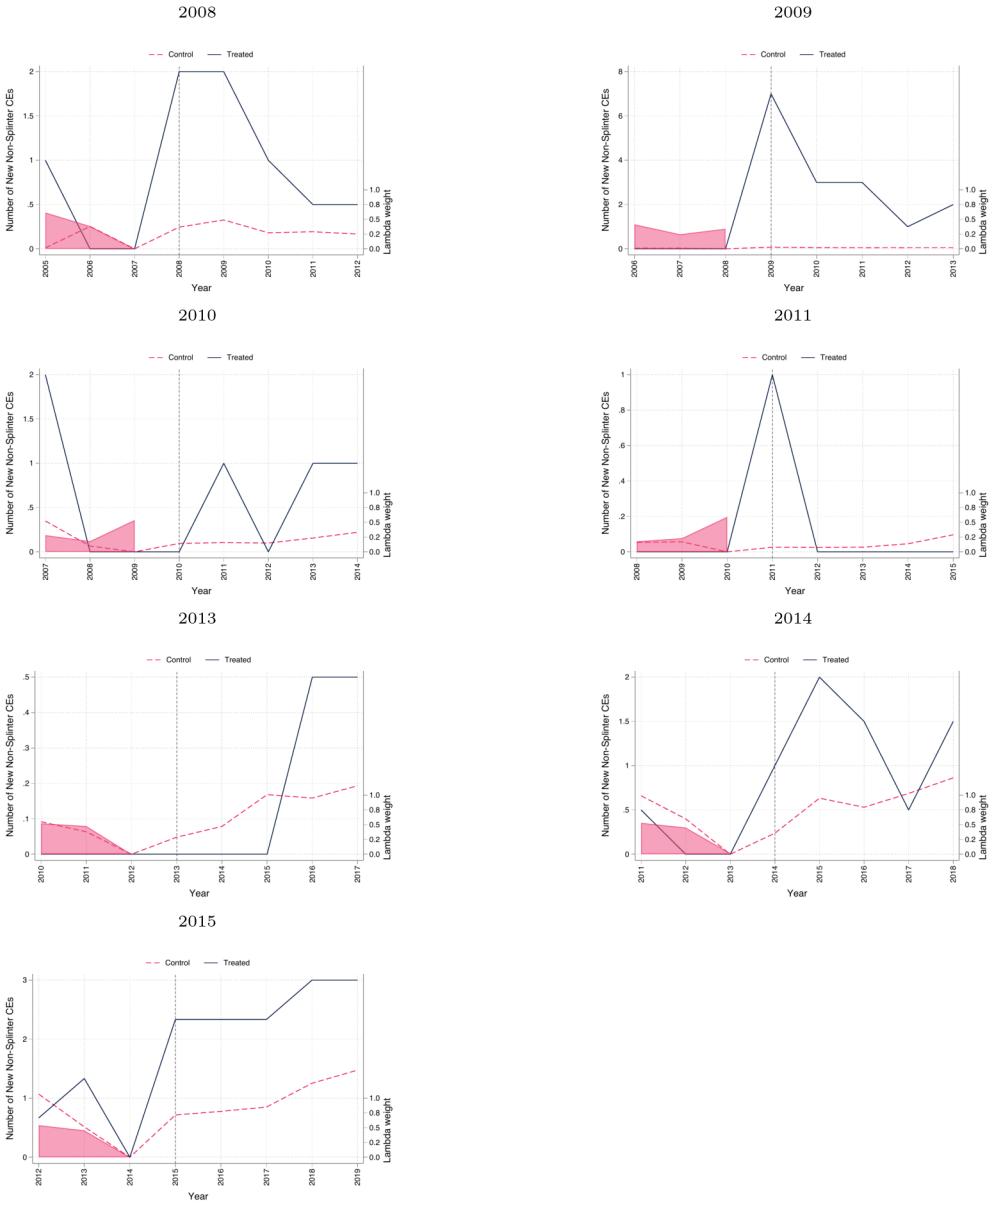

Notes: This figure plots, for each treatment date, the average number of non-new splinter criminal enterprises that were not present in the year before treatment for municipalities with kingpin captures and deaths in year t, along with the analogous ω weighted average for the synthetic control unit. The λ weights used to average pretreatment time periods are at the bottom of the graph. The control set X^{e}_{mt}, discussed in Section 3.4, includes interactions between year indicator variables and each of the following: The suitability of municipality m for opium poppy cultivation, the suitability of municipality m for cannabis cultivation, and an indicator of whether municipality m contains a Pacific port with shipping container service.

Appendix Figure A13. Number of New Splinters of CEs That Were Present in Year Before Treatment Over Time for Municipalities with Kingpin Captures and Synthetic Control Unit, by Treatment Date

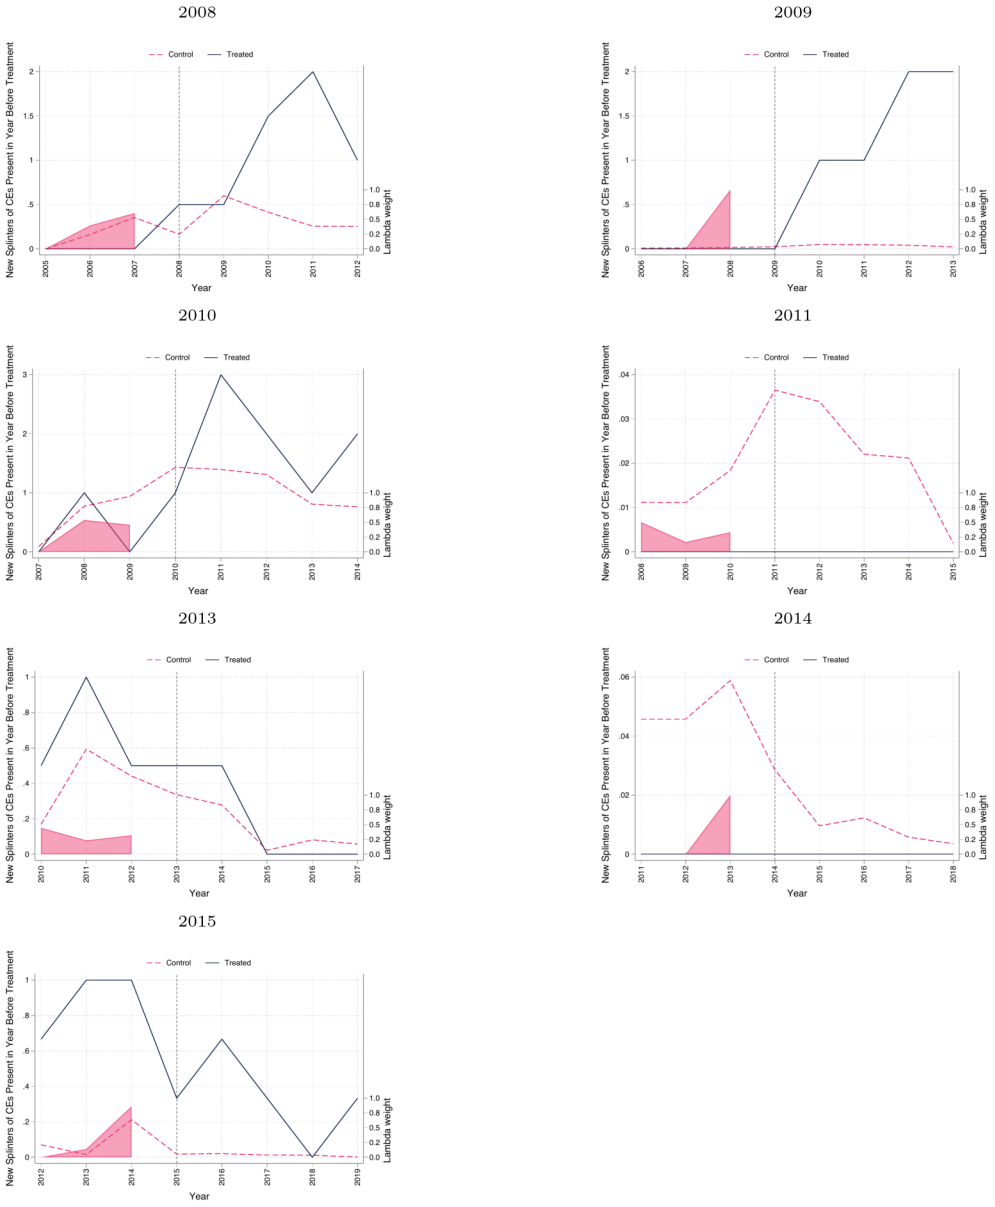

Notes: This figure plots, for each treatment date, the average number of new splinters of criminal enterprises that were present in the year before treatment for municipalities with kingpin captures and deaths in year t, along with the analogous ω weighted average for the synthetic control unit. The λ weights used to average pretreatment time periods are at the bottom of the graph. The control set X^e_{mt}, discussed in Section 3.4, includes interactions between year indicator variables and each of the following: The suitability of municipality m for opium poppy cultivation, the suitability of municipality m for cannabis cultivation, and an indicator of whether municipality m contains a Pacific port with shipping container service.

Appendix Figure A14. Number of New Splinters of CEs That Were Not Present in Year Before Treatment Over Time for Municipalities with Kingpin Captures and Synthetic Control Unit, by Treatment Date

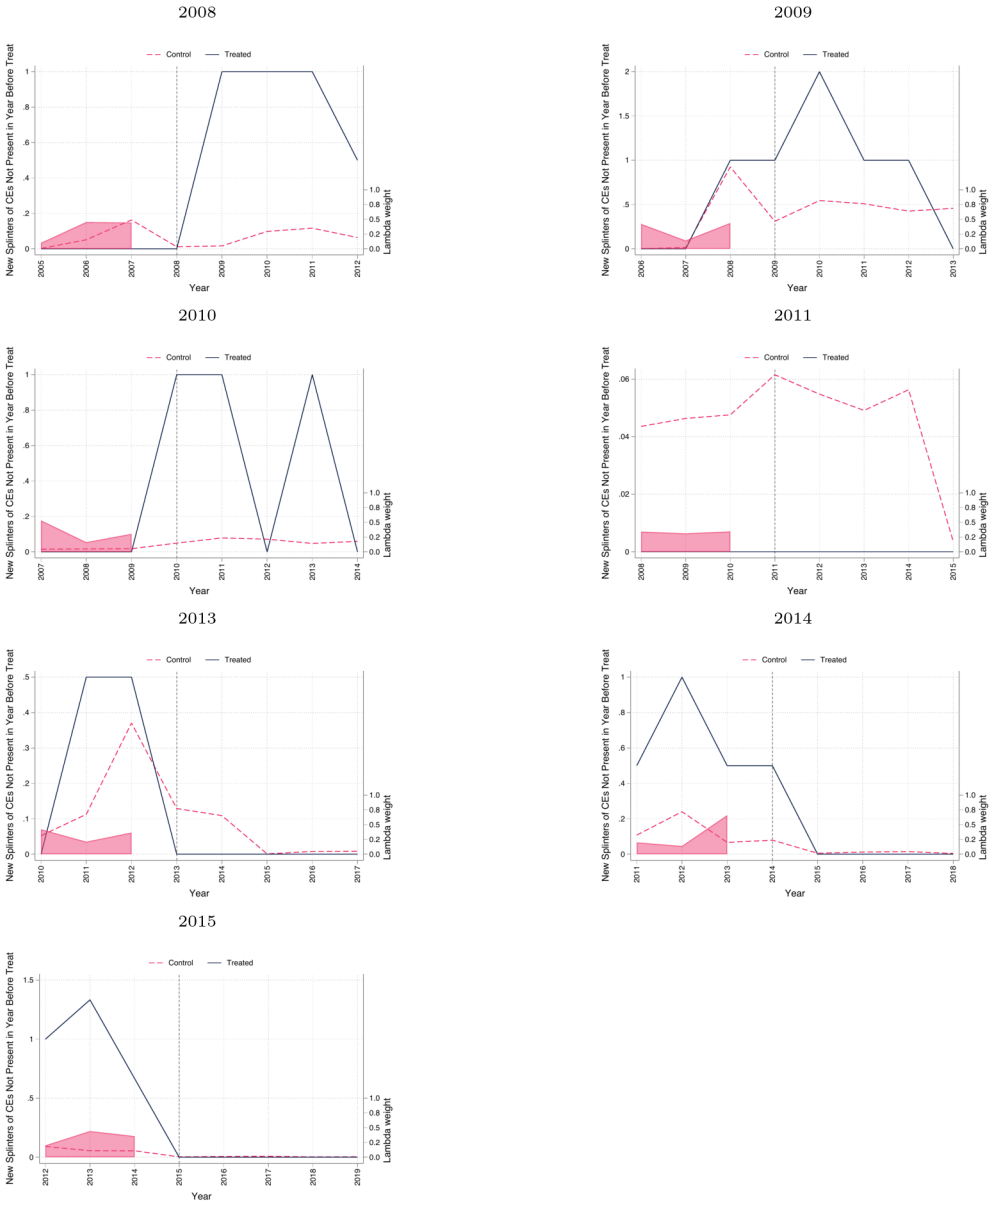

Notes: This figure plots, for each treatment date, the average number of new splinters of criminal enterprises that were not present in the year before treatment for municipalities with kingpin captures and deaths in year t, along with the analogous ω weighted average for the synthetic control unit. The λ weights used to average pretreatment time periods are at the bottom of the graph. The control set X^{e}_{mt}, discussed in Section 3.4, includes interactions between year indicator variables and each of the following: The suitability of municipality m for opium poppy cultivation, the suitability of municipality m for cannabis cultivation, and an indicator of whether municipality m contains a Pacific port with shipping container service.

Appendix Figure A15. Number of Homicides Per 100,000 Population Over Time for Municipalities with Kingpin Captures and Synthetic Control Unit, by Treatment Date

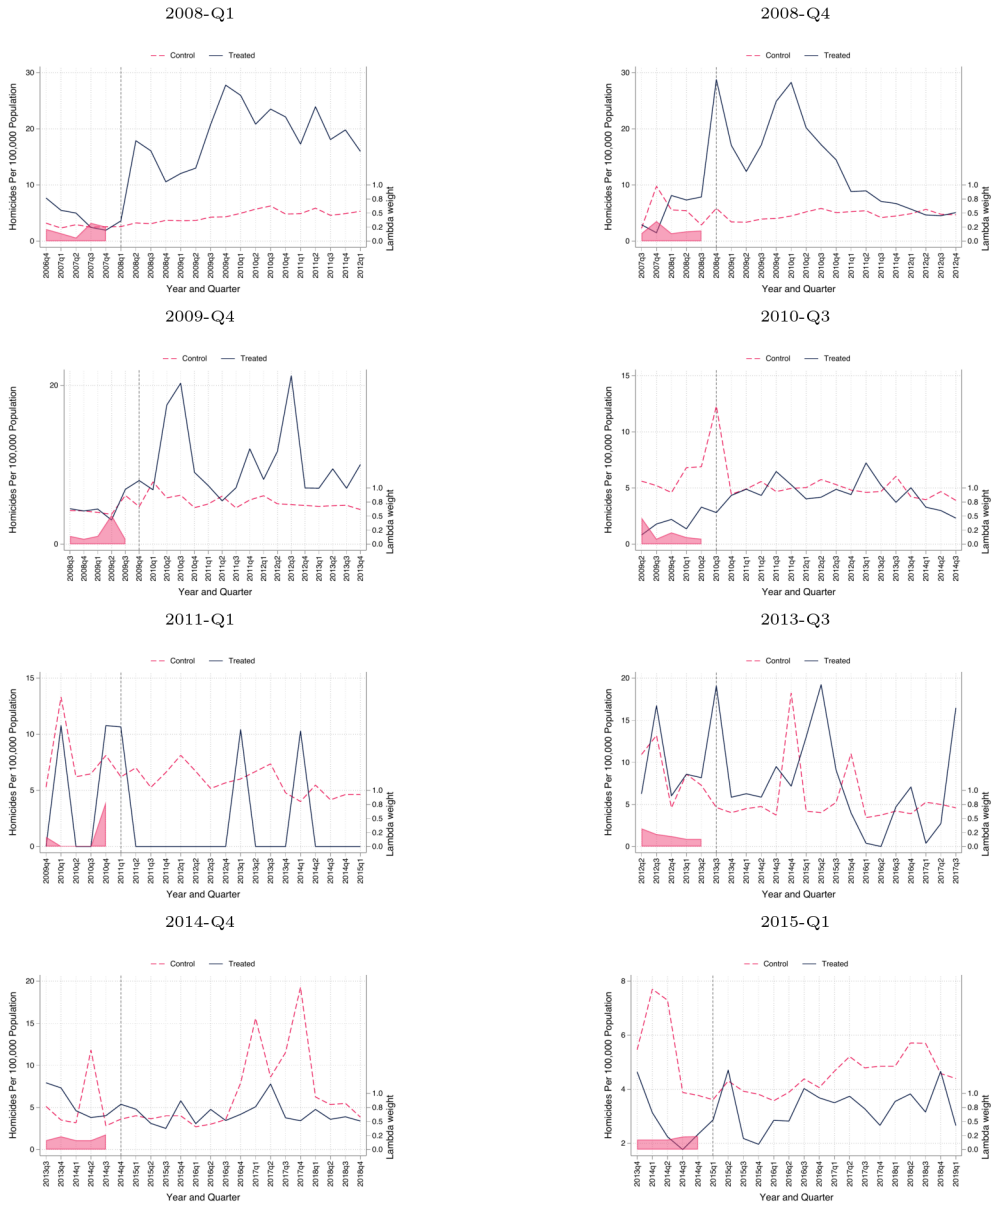

Notes: This figure plots, for each treatment date, the average number of homicides per 100,000 population for municipalities with kingpin captures and deaths in quarter t, along with the analogous ω weighted average for the synthetic control unit. The λ weights used to average pretreatment time periods are at the bottom of the graph. The control set X^{e}_{mt}, discussed in Section 3.4, includes interactions between year indicator variables and each of the following: The suitability of municipality m for opium poppy cultivation, the suitability of municipality m for cannabis cultivation, and an indicator of whether municipality m contains a Pacific port with shipping container service.

Appendix Figure A16. Number of Homicides of Males Aged 15-39 Per 100,000 Population Over Time for Municipalities with Kingpin Captures and Synthetic Control Unit, by Treatment Date

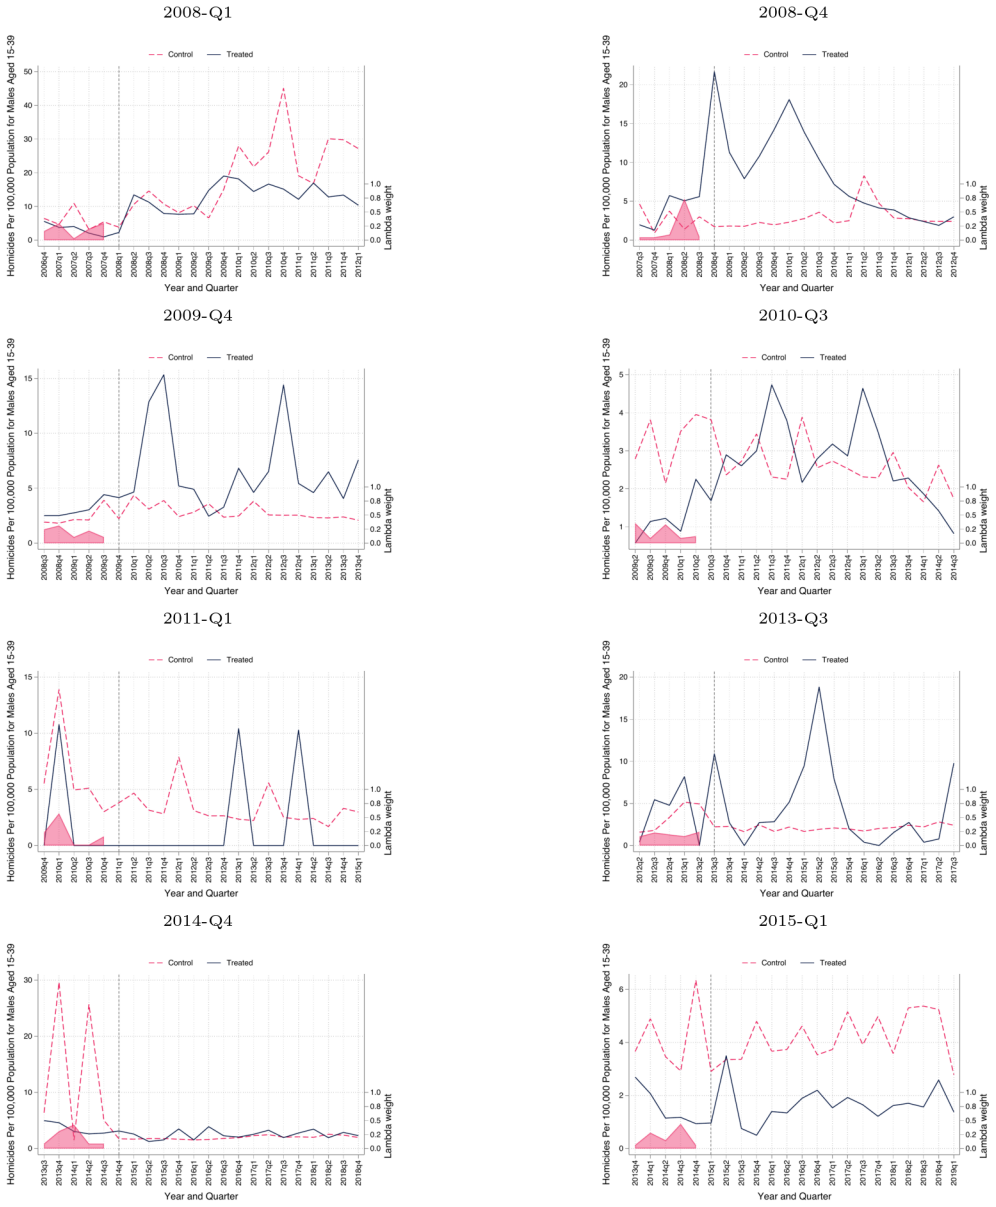

Notes: This figure plots, for each treatment date, the average number of homicides of males aged 15-39 per 100,000 population for municipalities with kingpin captures and deaths in quarter t, along with the analogous ω weighted average for the synthetic control unit. The λ weights used to average pretreatment time periods are at the bottom of the graph. The control set X^{e}_{mt}, discussed in Section 3.4, includes interactions between year indicator variables and each of the following: The suitability of municipality m for opium poppy cultivation, the suitability of municipality m for cannabis cultivation, and an indicator of whether municipality m contains a Pacific port with shipping container service.

Appendix Figure A17. Number of Homicides (Excluding Males Aged 15-39) Per 100,000 Population Over Time for Municipalities with Kingpin Captures and Synthetic Control Unit, by Treatment Date

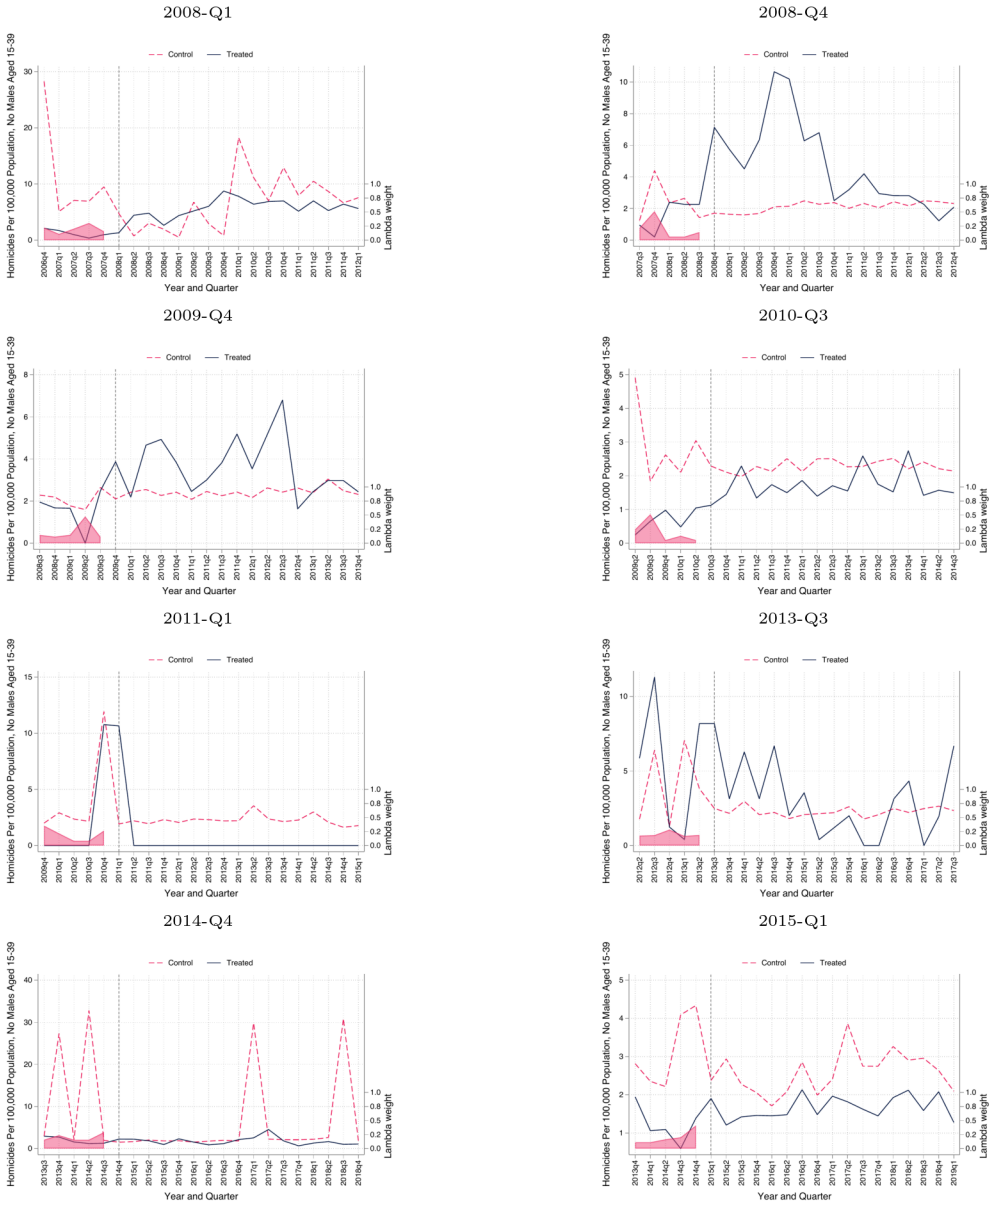

Notes: This figure plots, for each treatment date, the average number of homicides (excluding males aged 15-39) per 100,000 population for municipalities with kingpin captures and deaths in quarter t, along with the analogous ω weighted average for the synthetic control unit. The λ weights used to average pretreatment time periods are at the bottom of the graph. The control set X^e_{mt}, discussed in Section 3.4, includes interactions between year indicator variables and each of the following: The suitability of municipality m for opium poppy cultivation, the suitability of municipality m for cannabis cultivation, and an indicator of whether municipality m contains a Pacific port with shipping container service.

Appendix Figure A18. Distribution of Synthetic Control Weights by Treatment Cohort for Outcome: Number of Criminal Enterprises

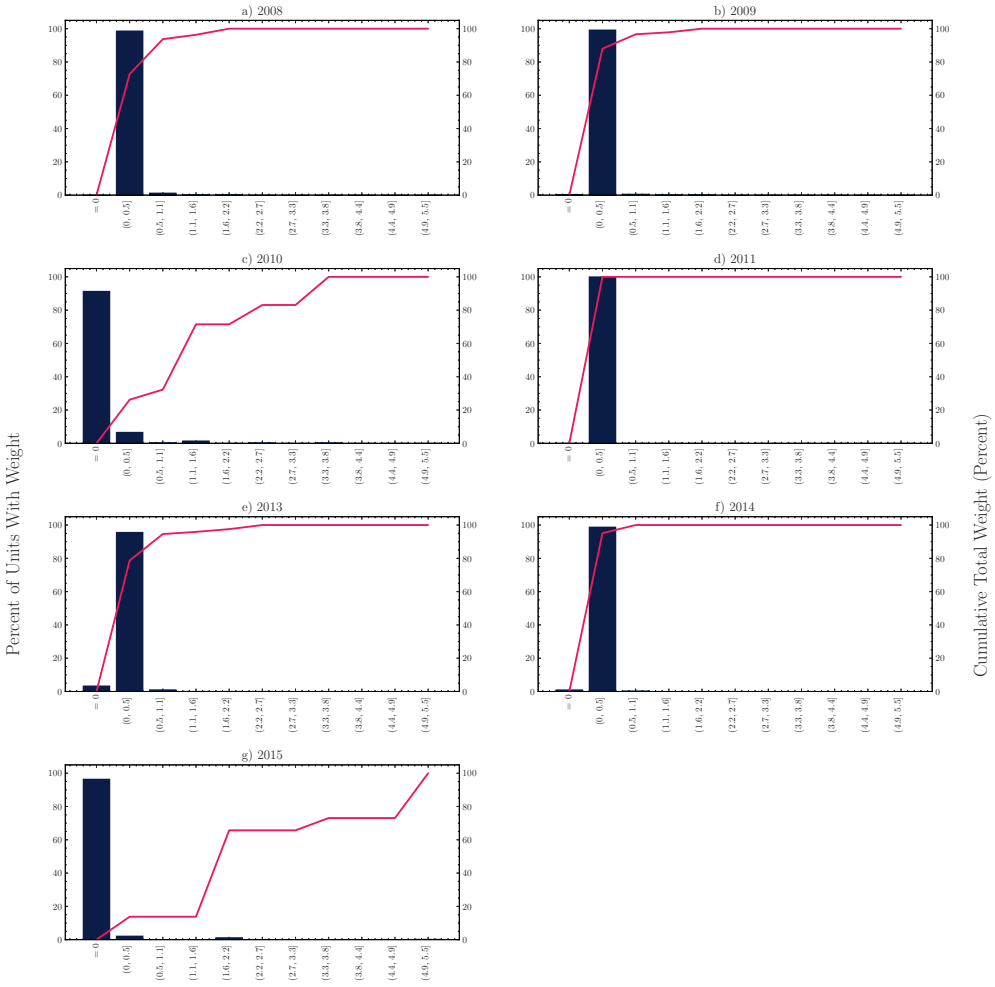

Notes: This figure plots the distribution of synthetic control ω weights for the synthetic difference-in-differences estimator for the effect of kingpin captures on the number of criminal enterprises by treatment cohort. Results are plotted for each cohort, which are composed of all municipalities that had a kingpin capture or death during year t and the never-treated control units used to compute the synthetic control for each cohort. The bars display the percentage of municipality units whose weights fall within specific bins—illustrating how the overall weight (which sums to 100% for each cohort) is distributed across municipalities. The pink line traces the cumulative contribution of these weights, showing the percentage of the total synthetic control constructed up to each bin. Bins are determined by allocating a zero-weight bin, and then equally dividing the remaining weight range (from above zero up to the maximum weight) into equal-sized intervals.

Appendix Figure A19. Distribution of Synthetic Control Weights by Treatment Cohort for Outcome: Number of Non-New Splinter CEs Present in Year Before Treatment

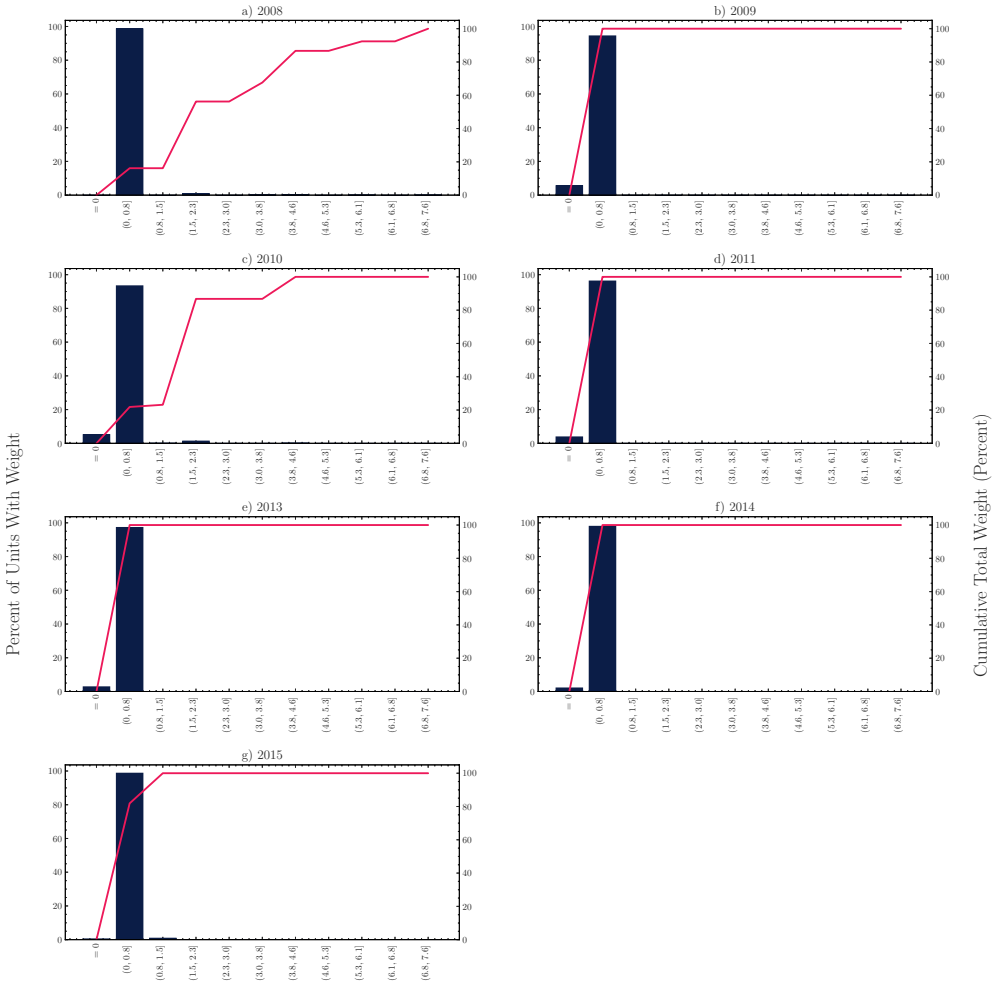

Notes: This figure plots the distribution of synthetic control ω weights for the synthetic difference-in-differences estimator for the effect of kingpin captures on the non-new splinter criminal enterprises that were present the year before treatment in municipality m in year t. Results are plotted for each cohort, which are composed of all municipalities that had a kingpin capture or death during year t and the never-treated control units used to compute the synthetic control for each cohort. The bars display the percentage of municipality units whose weights fall within specific bins—illustrating how the overall weight (which sums to 100% for each cohort) is distributed across municipalities. The pink line traces the cumulative contribution of these weights, showing the percentage of the total synthetic control constructed up to each bin. Bins are determined by allocating a zero-weight bin, and then equally dividing the remaining weight range (from above zero up to the maximum weight) into equal-sized intervals.

Appendix Figure A20. Distribution of Synthetic Control Weights by Treatment Cohort for Outcome: Number of Non-New Splinter CEs Not Present in Year Before Treatment

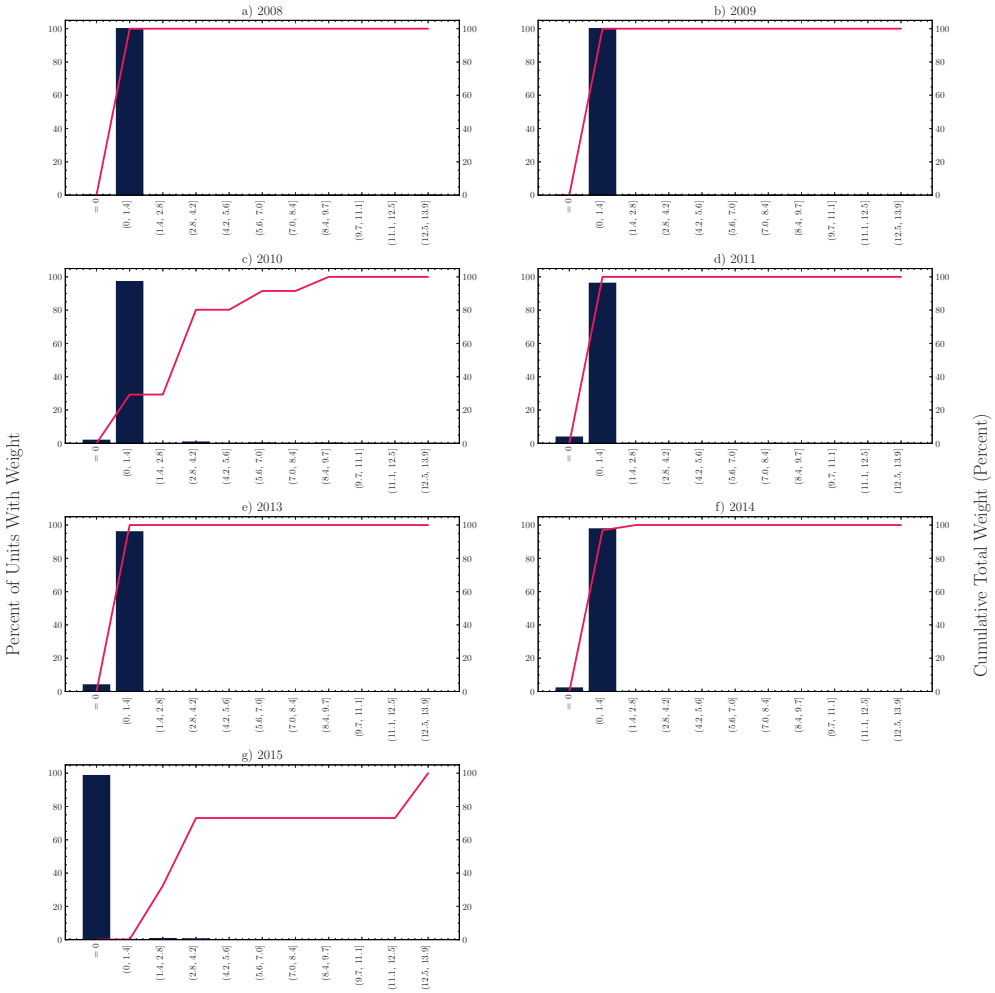

Notes: This figure plots the distribution of synthetic control ω weights for the synthetic difference-in-differences estimator for the effect of kingpin captures on the non-new splinter criminal enterprises that were not present the year before treatment in municipality m in year t. Results are plotted for each cohort, which are composed of all municipalities that had a kingpin capture or death during year t and the never-treated control units used to compute the synthetic control for each cohort. The bars display the percentage of municipality units whose weights fall within specific bins—illustrating how the overall weight (which sums to 100% for each cohort) is distributed across municipalities. The pink line traces the cumulative contribution of these weights, showing the percentage of the total synthetic control constructed up to each bin. Bins are determined by allocating a zero-weight bin, and then equally dividing the remaining weight range (from above zero up to the maximum weight) into equal-sized intervals.

Appendix Figure A21. Distribution of Synthetic Control Weights by Treatment Cohort for Outcome: Number of New Splinters of CEs Present in Year Before Treatment

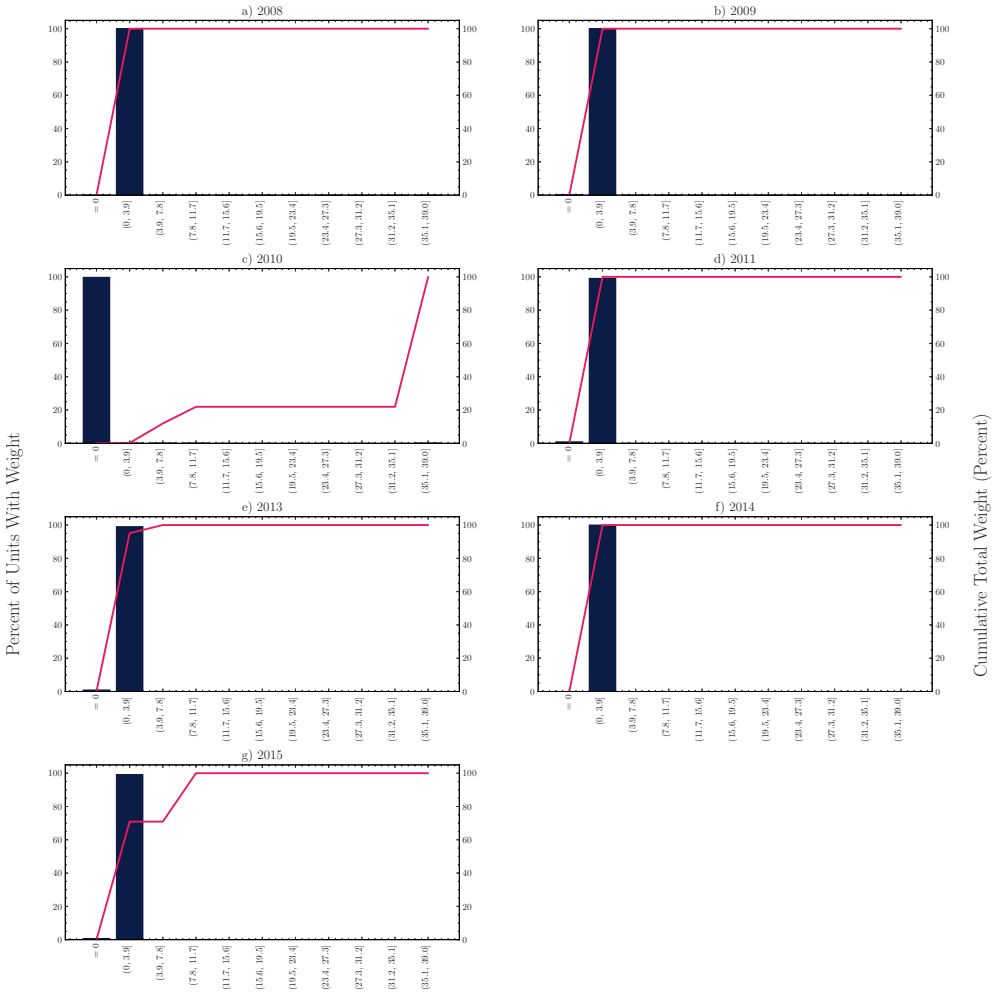

Notes: This figure plots the distribution of synthetic control ω weights for the synthetic difference-in-differences estimator for the effect of kingpin captures on the new splinter criminal enterprises of CEs that were present the year before treatment in municipality m in year t. Results are plotted for each cohort, which are composed of all municipalities that had a kingpin capture or death during year t and the never-treated control units used to compute the synthetic control for each cohort. The bars display the percentage of municipality units whose weights fall within specific bins—illustrating how the overall weight (which sums to 100% for each cohort) is distributed across municipalities. The pink line traces the cumulative contribution of these weights, showing the percentage of the total synthetic control constructed up to each bin. Bins are determined by allocating a zero-weight bin, and then equally dividing the remaining weight range (from above zero up to the maximum weight) into equal-sized intervals.

Appendix Figure A22. Distribution of Synthetic Control Weights by Treatment Cohort for Outcome: Number of New Splinters of CEs Not Present in Year Before Treatment

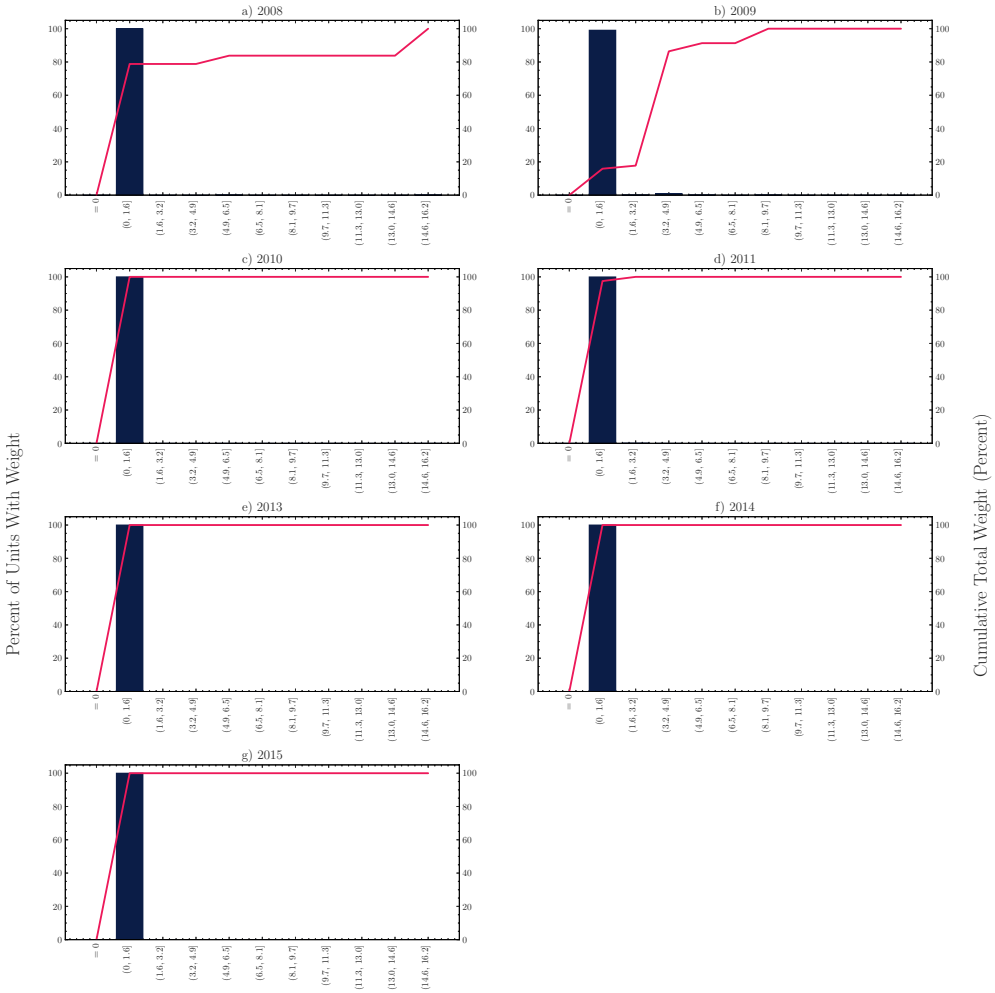

Notes: This figure plots the distribution of synthetic control ω weights for the synthetic difference-in-differences estimator for the effect of kingpin captures on the new splinter criminal enterprises of CEs that were not present the year before treatment in municipality m in year t. Results are plotted for each cohort, which are composed of all municipalities that had a kingpin capture or death during year t and the never-treated control units used to compute the synthetic control for each cohort. The bars display the percentage of municipality units whose weights fall within specific bins—illustrating how the overall weight (which sums to 100% for each cohort) is distributed across municipalities. The pink line traces the cumulative contribution of these weights, showing the percentage of the total synthetic control constructed up to each bin. Bins are determined by allocating a zero-weight bin, and then equally dividing the remaining weight range (from above zero up to the maximum weight) into equal-sized intervals.

Appendix Figure A23. Distribution of Synthetic Control Weights by Treatment Cohort for Outcome: Homicides Per 100,000 Population

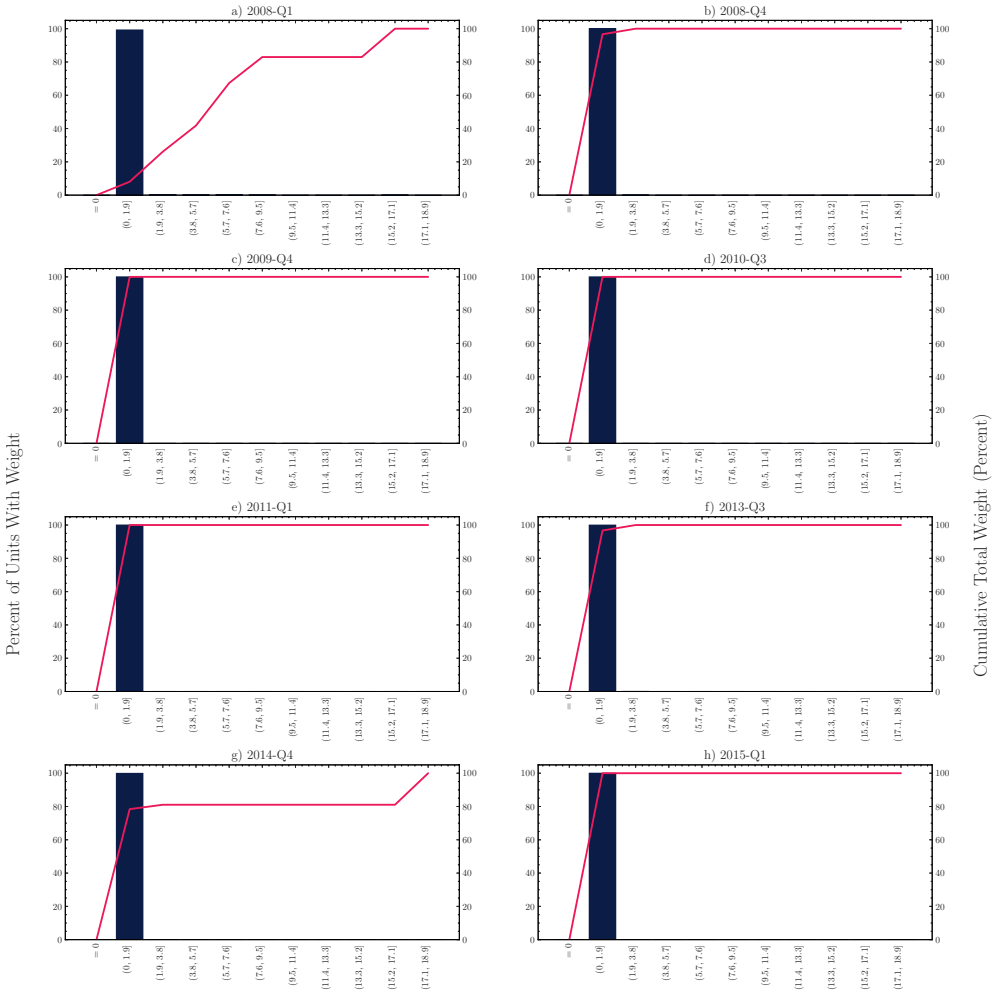

Notes: This figure plots the distribution of synthetic control ω weights for the synthetic difference-in-differences estimator for the effect of kingpin captures on the number of homicides per 100,000 population in municipality m in quarter t. Results are plotted for each cohort, which are composed of all municipalities that had a kingpin capture or death during quarter t and the never-treated control units used to compute the synthetic control for each cohort. The bars display the percentage of municipality units whose weights fall within specific bins—illustrating how the overall weight (which sums to 100% for each cohort) is distributed across municipalities. The pink line traces the cumulative contribution of these weights, showing the percentage of the total synthetic control constructed up to each bin. Bins are determined by allocating a zero-weight bin, and then equally dividing the remaining weight range (from above zero up to the maximum weight) into equal-sized intervals.

Appendix Figure A24. Distribution of Synthetic Control Weights by Treatment Cohort for Outcome: Homicides Per 100,000 Population (Males Aged 15-39)

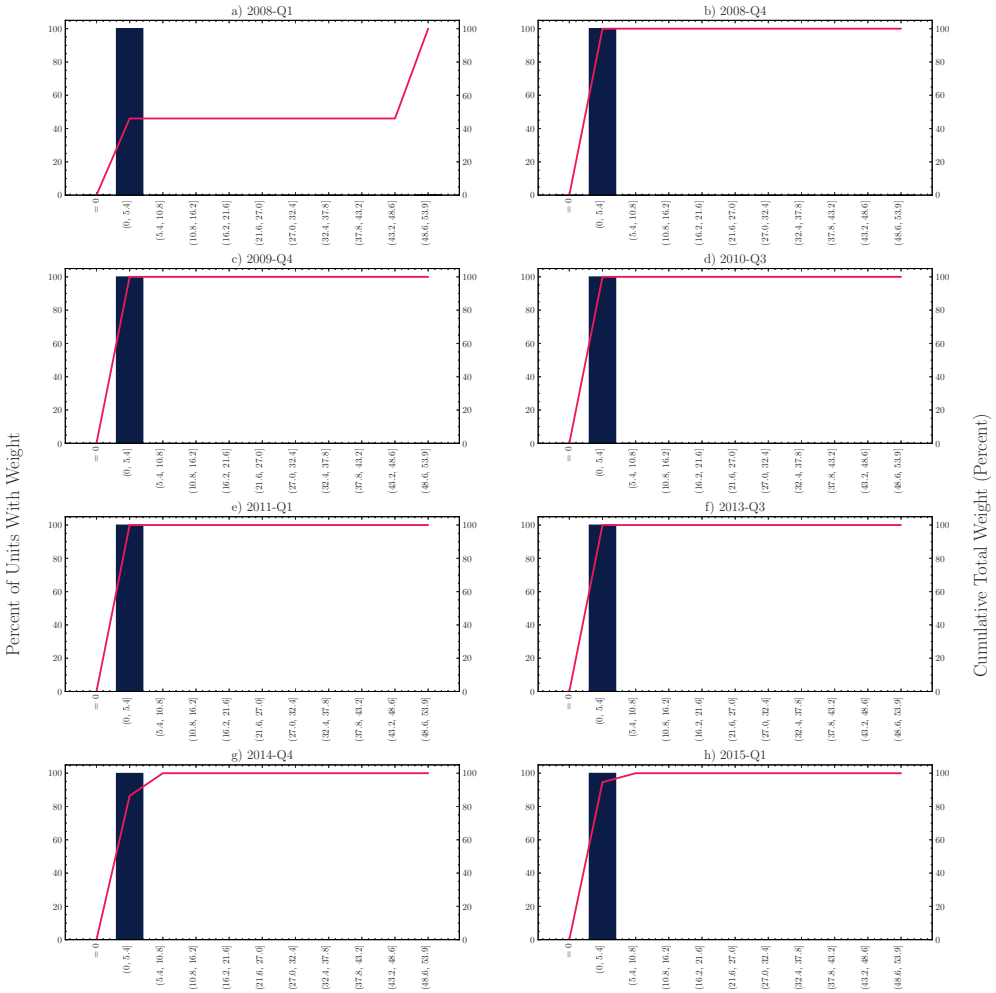

Notes: This figure plots the distribution of synthetic control ω weights for the synthetic difference-in-differences estimator for the effect of kingpin captures on the number of homicides of males aged 15-39 per 100,000 population in municipality m in quarter t. Results are plotted for each cohort, which are composed of all municipalities that had a kingpin capture or death during quarter t and the never-treated control units used to compute the synthetic control for each cohort. The bars display the percentage of municipality units whose weights fall within specific bins—illustrating how the overall weight (which sums to 100% for each cohort) is distributed across municipalities. The pink line traces the cumulative contribution of these weights, showing the percentage of the total synthetic control constructed up to each bin. Bins are determined by allocating a zero-weight bin, and then equally dividing the remaining weight range (from above zero up to the maximum weight) into equal-sized intervals.

Appendix Figure A25. Distribution of Synthetic Control Weights by Treatment Cohort for Outcome: Homicides Per 100,000 Population (Excluding Males Aged 15-39)

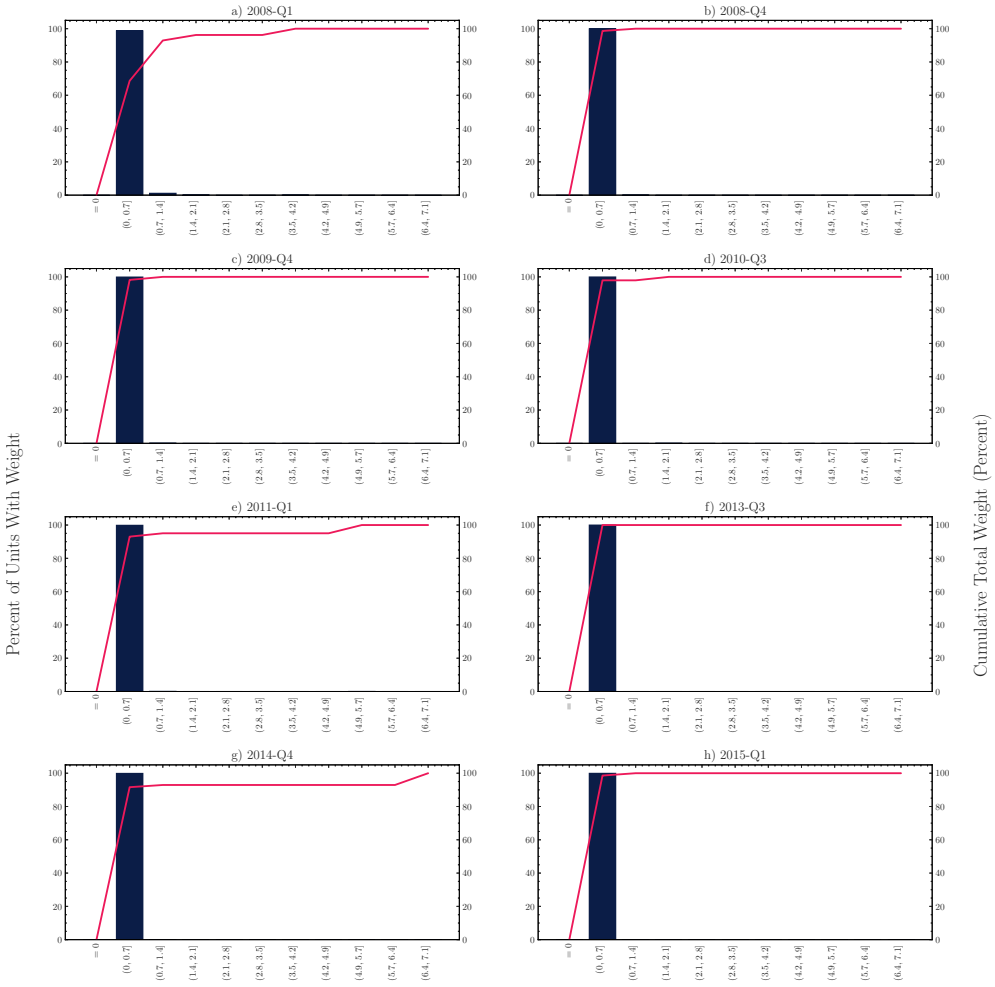

Notes: This figure plots the distribution of synthetic control ω weights for the synthetic difference-in-differences estimator for the effect of kingpin captures on the number of homicides of all people excluding males aged 15-39 per 100,000 population in municipality m in quarter t. Results are plotted for each cohort, which are composed of all municipalities that had a kingpin capture or death during quarter t and the never-treated control units used to compute the synthetic control for each cohort. The bars display the percentage of municipality units whose weights fall within specific bins—illustrating how the overall weight (which sums to 100% for each cohort) is distributed across municipalities. The pink line traces the cumulative contribution of these weights, showing the percentage of the total synthetic control constructed up to each bin. Bins are determined by allocating a zero-weight bin, and then equally dividing the remaining weight range (from above zero up to the maximum weight) into equal-sized intervals.

Appendix Figure A26. Examples of Narcomantas

(a) Narcomanta 1

(b) Narcomanta 2

Notes: These images show examples of *narcomantas*, large banners displayed in public spaces by criminal organizations to issue threats, assert dominance, or communicate with rivals, authorities, and the public. The banner in Image (a) reads: 'ATTENTION — Due to the incessant misinformation from some media outlets and the government's obvious failure to investigate and prosecute the real culprits of this epidemic, in Sinaloa, the sale, manufacture, transportation or any type of business involving the substance known as fentanyl is strictly prohibited, including the sale of chemicals for its production. We have never been and will never be related to that business. Face the consequences. Attn: Chapitos.' The banner in Image (b) reads: 'You're next Mayo, Tijuana belongs to the Arellanos.'

Appendix Figure A27. Recall With Mapping Criminal Organizations Project, By Group Prominence

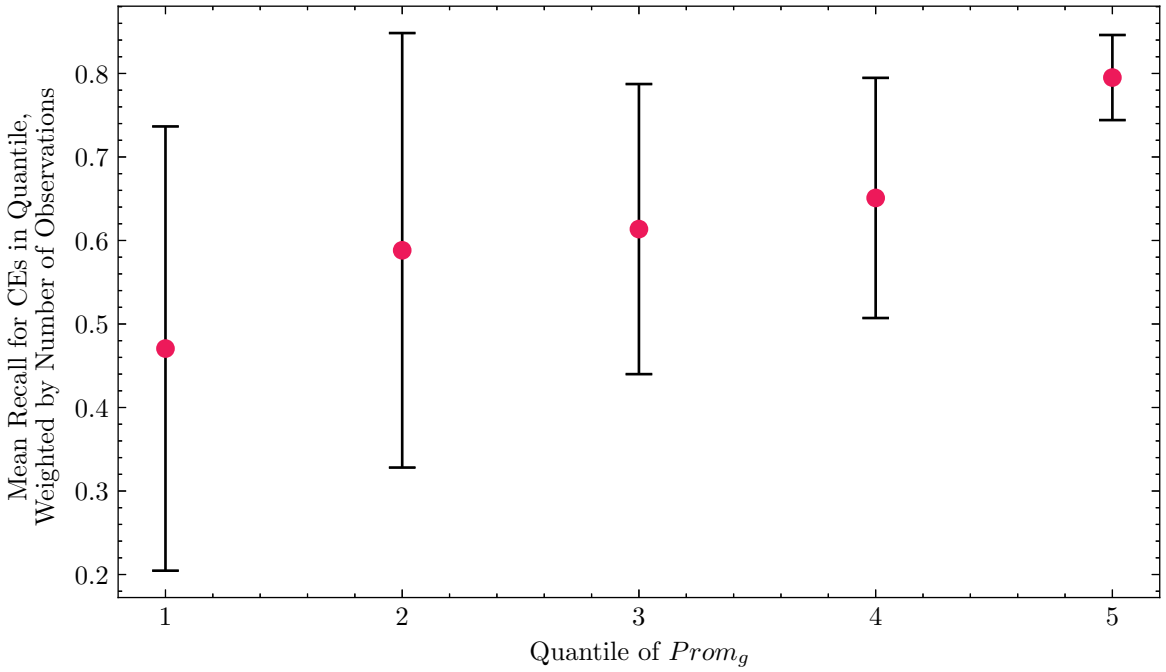

Notes: This figure shows the weighted mean recall across quintiles of CE prominence. Each CE's prominence is calculated based on the proportion of state-year combinations where it is present, adjusted for the number of observations, as defined in Equation 12. Recall measures the share of cases where a CE is predicted to be present in a state in a given year where it is actually present, using the MCO data as ground truth. Recall is aggregated to a quintile level average using the number of observations used to compute the recall for the quintile as weights.

Appendix Figure A28. Recall of CE Panel with Mapping Criminal Organizations Project, by Year

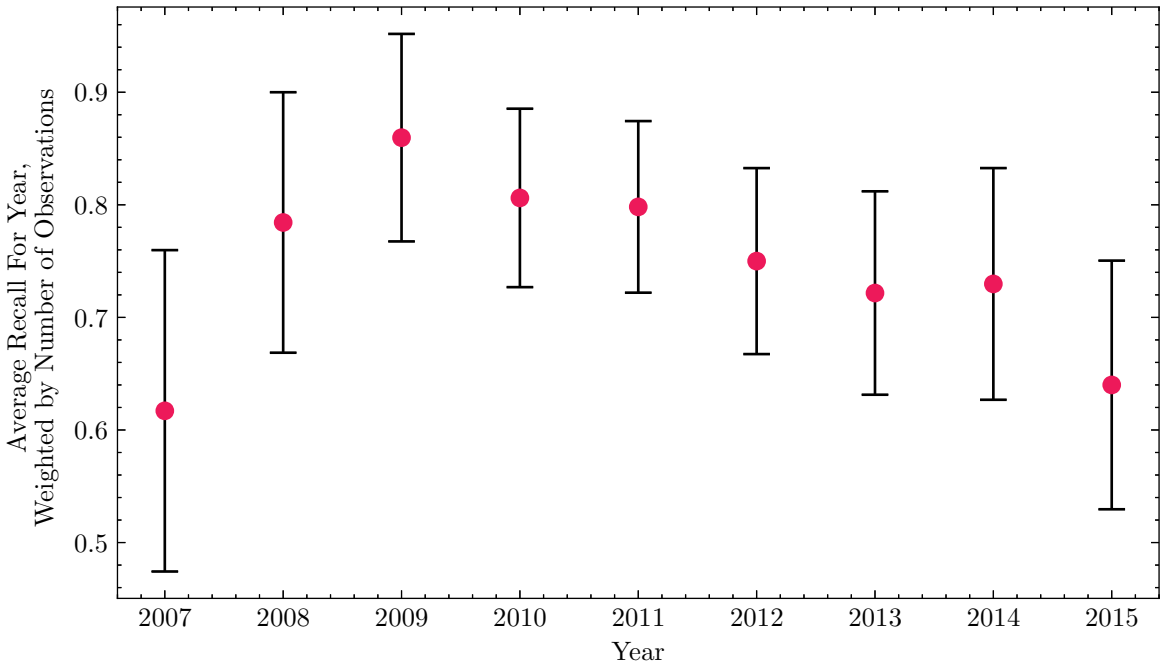

Notes: This figure shows the weighted mean for all criminal enterprises (CEs) over time, with 95% confidence intervals based on weighted standard errors. Recall measures the share of cases where a CE is predicted to be present in a state in a given year where it is actually present, using the MCO data as ground truth. Precision measures the share of cases where a CE is predicted to be present and is actually present. Both recall and precision are weighted by the number of observations for each CE-year combination.

Appendix Figure A29. Cumulative Number of Sources and Total Articles in Database, by Year

Notes: This figure shows the cumulative number of unique sources and the total number of full-text Spanish-language articles (in millions) over time from 2005 to 2019 in the ProQuest Latin American Newsstream. The shaded region represents the event window (2008-2015) used for the event study analysis. The cumulative number of sources accounts for sources added and removed each year, while the total number of articles represents the total annual publication volume.

Appendix Figure A30. Recall of CE Panel with Mapping Criminal Organizations Project, by State Population

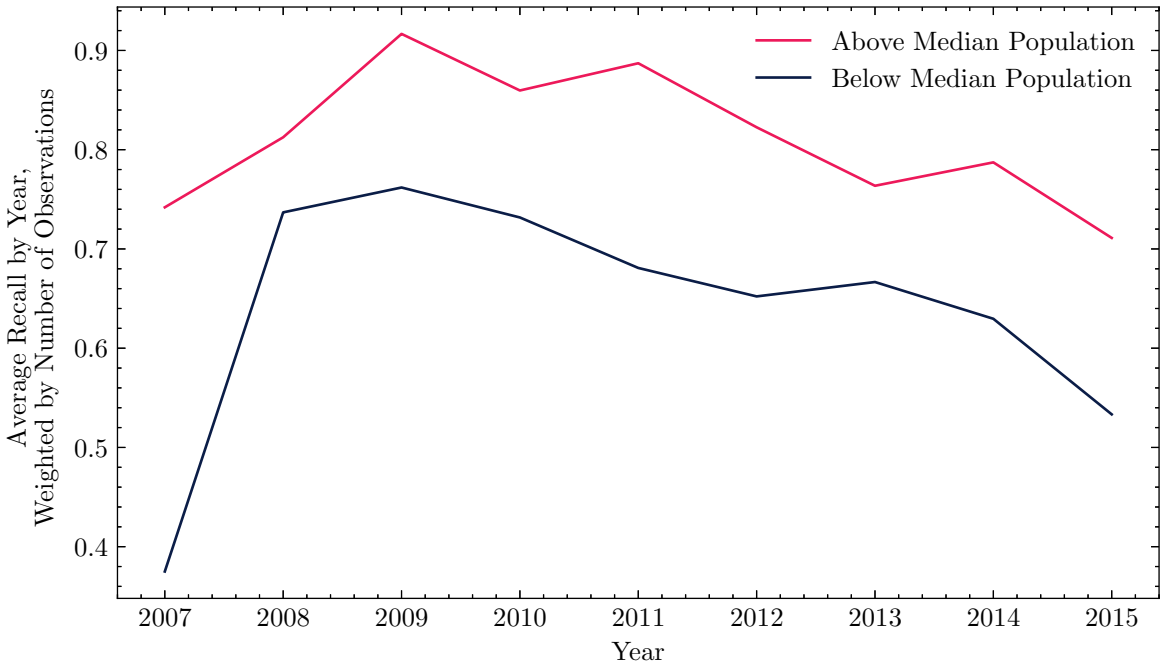

Notes: This figure shows the weighted mean for all criminal enterprises (CEs) over time, with 95% confidence intervals based on weighted standard errors. Recall measures the share of cases where a CE is predicted to be present in a state in a given year where it is actually present, using the MCO data as ground truth. Recall is weighted by the number of observations for each CE-year combination.

Appendix Figure A31. Effect of Kingpin Captures/Deaths on Non-Organized Crime Related Newspaper Coverage

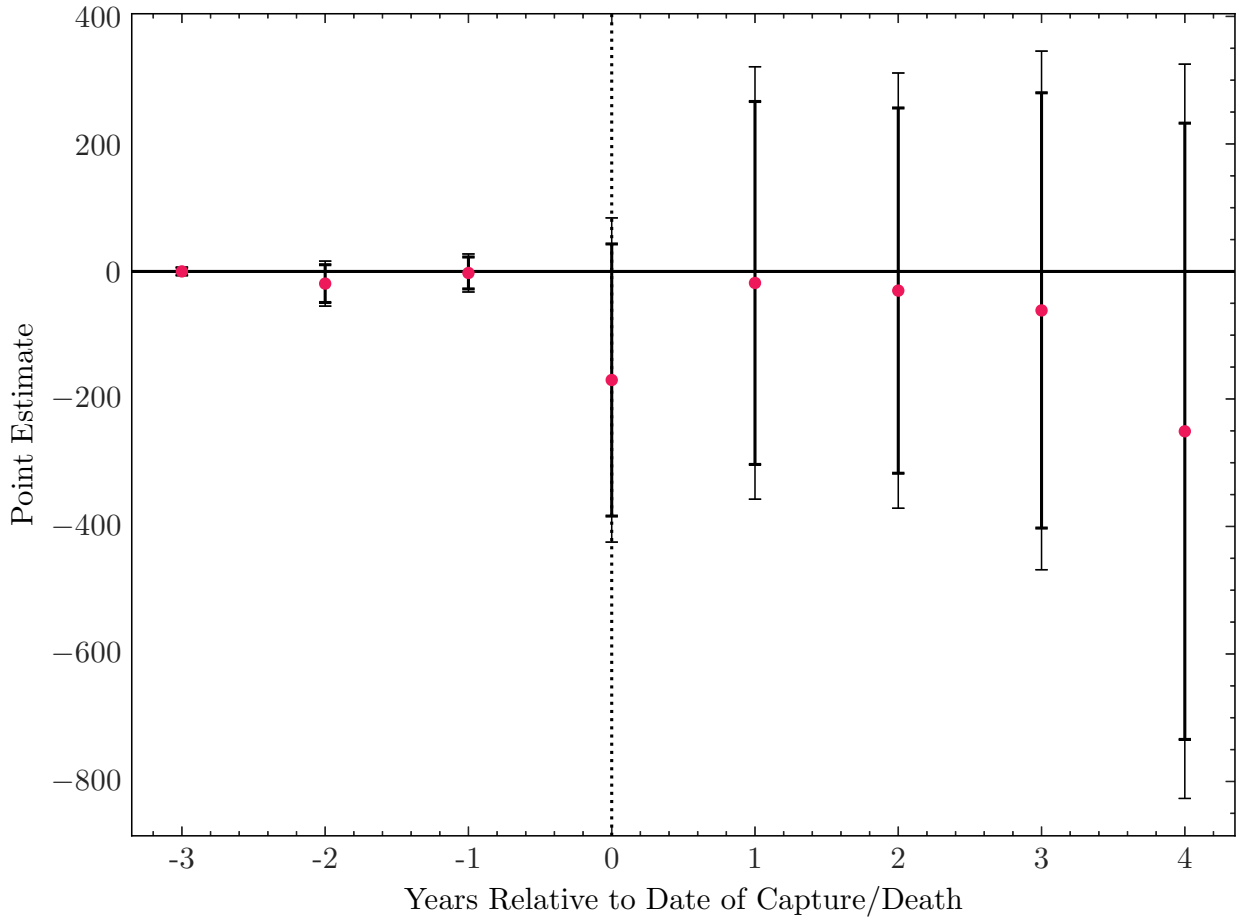

Notes: This figure plots estimates of the cohort-weighted average treatment effect on the treated $l \in \{-3, 4\}$ years relative to treatment, which is computed using the double difference event study estimator defined in Equation 10. The outcome is the number of unique non-organized crime related newspaper articles mentioning municipality m in year t, which is residualized on the set of time-variant controls X^e_{mt}. The panel for the outcome is built using the process described in Section 3.6. The cohort specific ATT is computed as the difference between $\phi^e_{tr,l}$ and $\phi^e_{co,l}$. For each cohort e, $\phi^e_{tr,l}$ is the average difference between the residualized outcome in treated municipality m in period l, relative to the time-weighted pre-period mean for the residualized outcome. $\phi^e_{co,l}$ is the unit-weighted difference between the residualized outcome l years relative to treatment and the time-weighted pre-period mean for the never-treated control units. Cohort-level estimates are aggregated using weights proportional to the number of treated units in each cohort, as defined in Equation 7. The control set X^e_{mt}, discussed in Section 3.4, includes interactions between year indicator variables and each of the following: The suitability of municipality m for opium poppy cultivation, the suitability of municipality m for cannabis cultivation, and an indicator of whether municipality m contains a Pacific port with shipping container service. The unit and time weights are estimated following Arkhangelsky et al. (2021) as defined in Appendix Section C.1. The data is derived from newspaper articles in the ProQuest Latin American Newsstream using the data processing pipeline outlined in Section 3.1. Standard errors are obtained through block bootstrapping with 100 replications as outlined in Clarke et al. (2023) with sampling occurring at the state level to account for spatial correlation in treatment assignment. The bold vertical bands represent 90% confidence intervals. The other vertical bands represent 95% confidence intervals.

Appendix Figure A32. Effect of Kingpin Captures/Deaths on
Number of National Guard and Federal Police Units Deployed to State

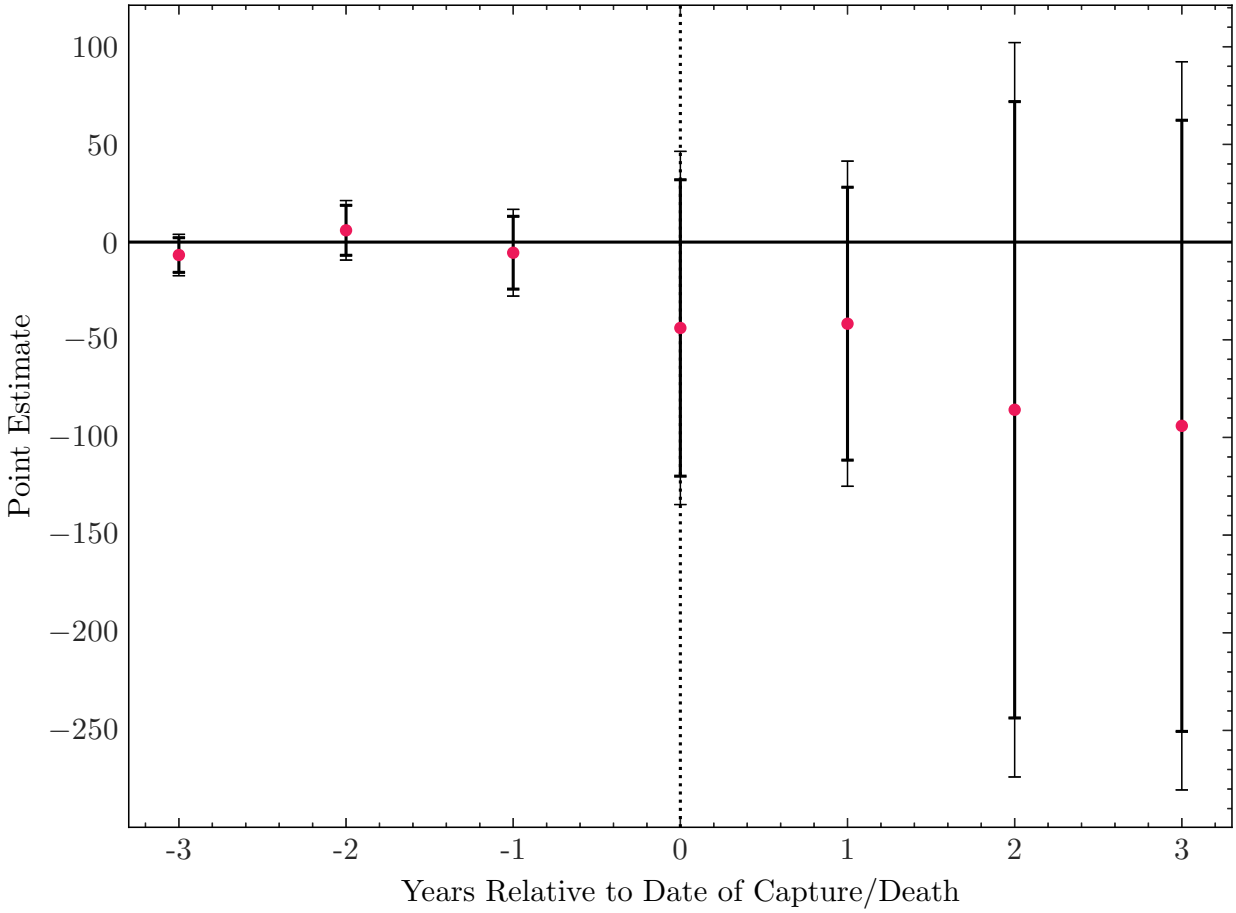

Notes: This figure plots estimates of the cohort-weighted average treatment effect on the treated $l \in \{-3, 3\}$ years relative to treatment, which is computed using the double difference event study estimator defined in Equation 10. The outcome is the number of National Guard and Federal Police units operating in state s in year t, which is residualized on the set of time-variant controls X^e_{st}. The cohort specific ATT is computed as the difference between $\phi^e_{tr,l}$ and $\phi^e_{co,l}$. For each cohort e, $\phi^e_{tr,l}$ is the average difference between the residualized outcome in treated state s in period l, relative to the time-weighted pre-period mean for the residualized outcome. $\phi^e_{co,l}$ is the unit-weighted difference between the residualized outcome l years relative to treatment and the time-weighted pre-period mean for the never-treated control units. Cohort-level estimates are aggregated using weights proportional to the number of treated states in each cohort, as defined in Equation 7. The control set X^e_{st}, discussed in Section 3.4, includes interactions between year indicator variables and each of the following: The suitability of state s for opium poppy cultivation, the suitability of state s for cannabis cultivation, and an indicator of whether municipality s contains a Pacific port with shipping container service. The unit and time weights are estimated following Arkhangelsky et al. (2021) as defined in Appendix Section C.1. Standard errors are obtained through block bootstrapping with 100 replications as outlined in Clarke et al. (2023) with sampling occurring at the state level to account for spatial correlation in treatment assignment. The bold vertical bands represent 90% confidence intervals. The other vertical bands represent 95% confidence intervals.

Appendix Figure A33. Robustness of Main Results to the Inclusion of All Kingpins

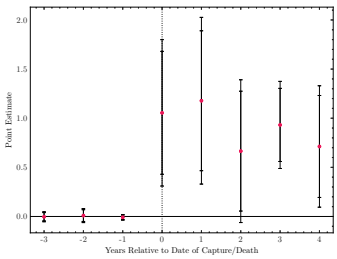

(a) Number of Criminal Enterprises

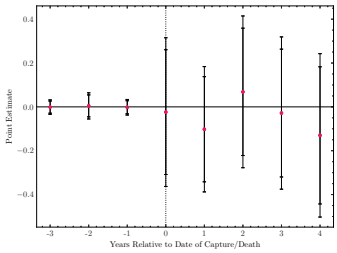

(b) Number of Non-New Splinter CEs Present in Year Before Treatment

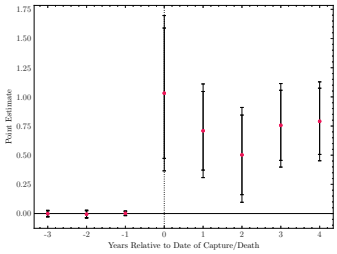

(c) Number of Non-New Splinter CEs Not Present in Year Before Treatment

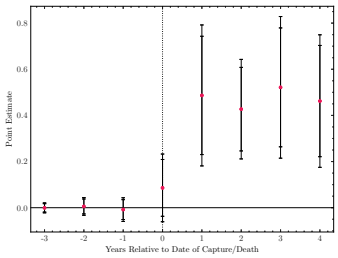

(d) Number of New Splinters of CEs Present in Year Before Treatment

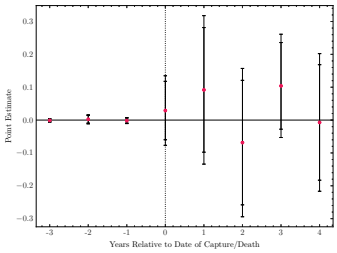

(e) Number of New Non Splinters of CEs Not Present in Year Before Treatment

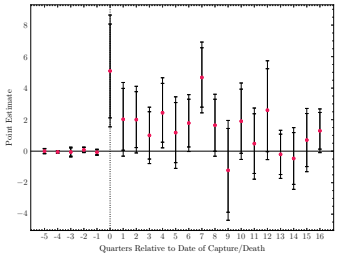

(f) Homicides Per 100,000 Population

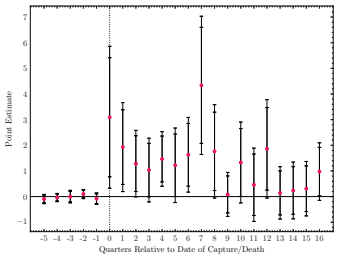

(g) Homicide Rate - Victim's Demographics Match Those of Typical CE Member

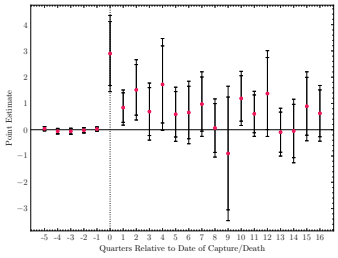

(h) Homicide Rate - Victim Unlikely CE Member

Notes: This figure shows the robustness of the main results to the inclusion of potentially kingpins that were excluded from the main analysis due to endogeneity concerns (discussed in Section 3.3). This figure plots estimates of the cohort-weighted average treatment effect on the treated for the 4 years following treatment along with placebo estimates, which are computed using the double difference event study estimator defined in Equation 10. The cohort specific ATT is computed as the difference between $\phi^e_{tr,l}$ and $\phi^e_{co,l}$. For each cohort e, $\phi^e_{tr,l}$ is the average difference between the residualized outcome in treated municipality m in period l, relative to the time-weighted pre-period mean for the residualized outcome. $\phi^e_{co,l}$ is the unit-weighted difference between the residualized outcome l years relative to treatment and the time-weighted pre-period mean for the never-treated control units. Cohort-level estimates are aggregated using weights proportional to the number of treated municipalities in each cohort, as defined in Equation 7. The control set X^e_{mt}, discussed in Section 3.4, includes interactions between year indicator variables and each of the following: The suitability of municipality m for opium poppy cultivation, the suitability of municipality m for cannabis cultivation, and an indicator of whether municipality s contains a Pacific port with shipping container service. The unit and time weights are estimated following Arkhangelsky et al. (2021) as defined in Appendix Section C.1. Standard errors are obtained through block bootstrapping with 100 replications as outlined in Clarke et al. (2023) with sampling occurring at the state level to account for spatial correlation in treatment assignment. The bold vertical bands represent 90% confidence intervals. The other vertical bands represent 95% confidence intervals.

Appendix Figure A34. Robustness of Main Results to Non-Imputation of CE Data

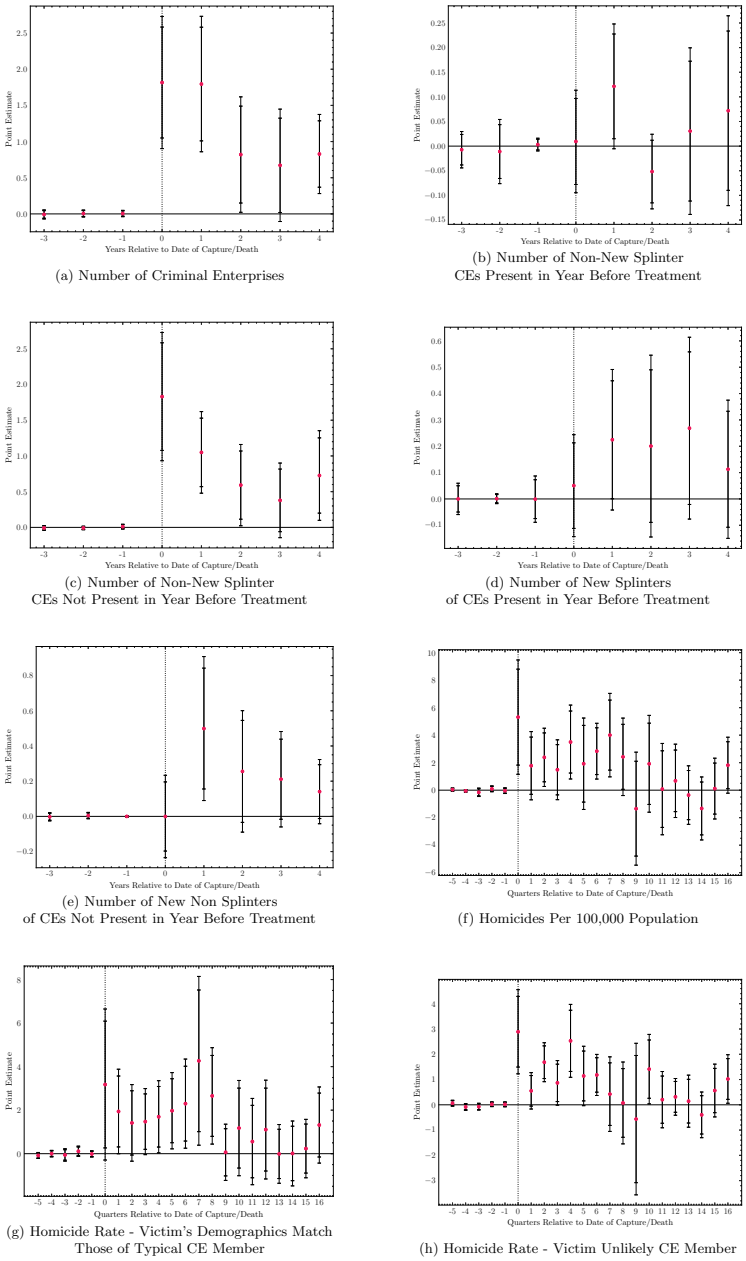

(a) Number of Criminal Enterprises

(b) Number of Non-New Splinter CEs Present in Year Before Treatment

(c) Number of Non-New Splinter CEs Not Present in Year Before Treatment

(d) Number of New Splinters of CEs Present in Year Before Treatment

(e) Number of New Non Splinters of CEs Not Present in Year Before Treatment

(f) Homicides Per 100,000 Population

(g) Homicide Rate - Victim's Demographics Match Those of Typical CE Member

(h) Homicide Rate - Victim Unlikely CE Member

Notes: This figure shows the robustness of the main results when CEs are not imputed. This figure plots estimates of the cohort-weighted average treatment effect on the treated for the 4 years following treatment along with placebo estimates, which are computed using the double difference event study estimator defined in Equation 10. The cohort specific ATT is computed as the difference between $\phi^e_{tr,l}$ and $\phi^e_{co,l}$. For each cohort e, $\phi^e_{tr,l}$ is the average difference between the residualized outcome in treated municipality m in period l, relative to the time-weighted pre-period mean for the residualized outcome. $\phi^e_{co,l}$ is the unit-weighted difference between the residualized outcome l years relative to treatment and the time-weighted pre-period mean for the never-treated control units. Cohort-level estimates are aggregated using weights proportional to the number of treated municipalities in each cohort, as defined in Equation 7. The control set X^e_{mt}, discussed in Section 3.4, includes interactions between year indicator variables and each of the following: The suitability of municipality m for opium poppy cultivation, the suitability of municipality m for cannabis cultivation, and an indicator of whether municipality s contains a Pacific port with shipping container service. The unit and time weights are estimated following Arkhangelsky et al. (2021) as defined in Appendix Section C.1. Standard errors are obtained through block bootstrapping with 100 replications as outlined in Clarke et al. (2023) with sampling occurring at the state level to account for spatial correlation in treatment assignment. The bold vertical bands represent 90% confidence intervals. The other vertical bands represent 95% confidence intervals.

Appendix Figure A35. Robustness of Main Results to Alternate Treatment Geography

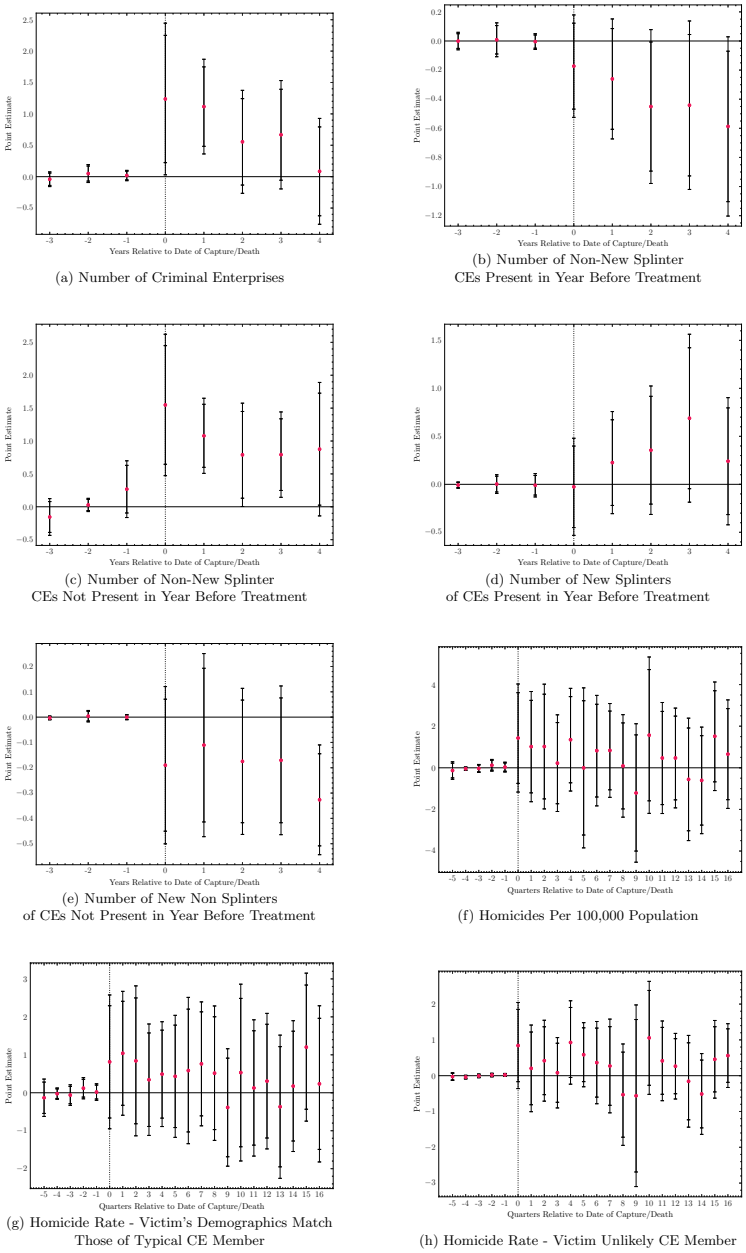

(a) Number of Criminal Enterprises

(b) Number of Non-New Splinter CEs Present in Year Before Treatment

(c) Number of Non-New Splinter CEs Not Present in Year Before Treatment

(d) Number of New Splinters of CEs Present in Year Before Treatment

(e) Number of New Non Splinters of CEs Not Present in Year Before Treatment

(f) Homicides Per 100,000 Population

(g) Homicide Rate - Victim's Demographics Match Those of Typical CE Member

(h) Homicide Rate - Victim Unlikely CE Member

Notes: This figure shows the robustness of the main results to an alternative definition of treated geography, as discussed in Section 5.5. Treated geographies now include outcomes aggregate across their immediate neighbors. This figure plots estimates of the cohort-weighted average treatment effect on the treated for the 4 years following treatment along with placebo estimates, which are computed using the double difference event study estimator defined in Equation 10. The cohort specific ATT is computed as the difference between $\phi^e_{tr,l}$ and $\phi^e_{co,l}$. For each cohort e, $\phi^e_{tr,l}$ is the average difference between the residualized outcome in treated municipality m in period l, relative to the time-weighted pre-period mean for the residualized outcome. $\phi^e_{co,l}$ is the unit-weighted difference between the residualized outcome l years relative to treatment and the time-weighted pre-period mean for the never-treated control units. Cohort-level estimates are aggregated using weights proportional to the number of treated municipalities in each cohort, as defined in Equation 7. The control set X^e_{mt}, discussed in Section 3.4, includes interactions between year indicator variables and each of the following: The suitability of municipality m for opium poppy cultivation, the suitability of municipality m for cannabis cultivation, and an indicator of whether municipality s contains a Pacific port with shipping container service. The unit and time weights are estimated following Arkhangelsky et al. (2021) as defined in Appendix Section C.1. Standard errors are obtained through block bootstrapping with 100 replications as outlined in Clarke et al. (2023) with sampling occurring at the state level to account for spatial correlation in treatment assignment. The bold vertical bands represent 90% confidence intervals. The other vertical bands represent 95% confidence intervals.

Appendix Table A1. Use of High Value Targeting in Major Drug Trafficking or Drug Producing Countries

Country	HVT Used	Evidence	Source(s)
Afghanistan	No	The DEA reports that attempts to extradite kingpins such as Mohammed Zia Salehi have been thwarted by the Afghani Government.	SIGAR
The Bahamas	Yes	Extradition of drug traffickers like Melvin Maycock Sr. involved in cocaine smuggling.	Tribune
Belize	No	There is no evidence of high-value targeting.	
Bolivia	No	There is no evidence of high-value targeting. Furthermore, Bolivia's President expelled the DEA from the country in 2008.	Pulitzer Center
China	No	Crackdowns on high ranking leaders of Chinese CEs is selective and rare, and depends on whether it aligns with government interests.	Brookings Institute
Colombia	Yes	As discussed in Section 2.3, Mexico's implementation of HVT was inspired by Colombia's use of the policy.	Lessing (2017)
Costa Rica	Yes	Capture of Alejandro Arias Monge, aka "Diablo", in a national operation.	Tico Times
Dominican Republic	Yes	Extradition of Julio De Los Santos-Bautista (Julito Kilo) for cocaine trafficking.	DR1
Ecuador	Yes	Extradition of Brayan Rodriguez Alcala, member of the Sinaloa Cartel.	ICE
El Salvador	Yes	Arrest of top MS-13 leaders, such as Armando Melgar Díaz.	InSight Crime
Guatemala	Yes	Extradition of Augusto Jean Carlo Castillo-Hernandez for cocaine trafficking.	ICE
Haiti	Yes	Extradition of Grégory Georges, aka "Ti Ketan", for heroin smuggling.	HaitiLibre
Honduras	Yes	Targeting of kingpin of the Montes Bobadilla drug trafficking clan.	ADN America
India	No	No clear evidence of high-value targeting for drug trafficking.	
Jamaica	Yes	Extradition of Christopher "Dudus" Coke, leader of the Shower Posse.	InSight Crime
Laos	No	There is no evidence of high-value targeting.	
Mexico	Yes	See Table 1 for examples of HVT in Mexico.	See Section 2.3.
Myanmar	No	There is no evidence of high-value targeting. Furthermore, there is high-level military involvement in drug trafficking.	The Diplomat
Nicaragua	No	There is no evidence of high-value targeting.	
Pakistan	Yes	There is no evidence of high-value targeting.	
Panama	Yes	Extradition of Piero Antonio Lubo-Barros for cocaine trafficking.	La Teja
Peru	Yes	Peruvian authorities captured Víctor Quispe Palomino, a top leader of the Shining Path.	DOJ
Venezuela	No	There is no evidence of high-value targeting. Furthermore, top-level government officials are involved in drug trafficking.	CSIS

Notes: This table lists, for each country on the Memorandum on Presidential Determination on Major Drug Transit or Major Illicit Drug Producing Countries for Fiscal Year 2024, whether or not there is evidence of the country's government using high value targeting against criminal enterprises. If there is, I list example(s) of this and the source where the example is discussed. This list of countries is published by the White House on an annual basis. Countries are placed on the list if there is a combination of geographic, commercial, and economic factors that allow drugs to be transited or produced, even if a government has engaged in robust and diligent narcotics control and law enforcement measures.

Appendix Table A2. Exogeneity of Kingpin Captures/Deaths

Leader Name	Included	Capture/Death Details	Source(s)
Alfredo Beltrán Leyva	Yes	Arrested in a short military raid that lasted 4 hours. While there are limited details on what led to the raid, a writer on military affairs, Javier Ibarrola, claimed that the capture appeared to be a nice piece of intelligence work by the federal government, which had said at the start of the year that its strategy had changed from patrolling streets in drug hot spots to more surgical strikes against top cartel operators."	La Journada The Seattle Times
Eduardo Arellano Félix	Yes	Details are scarce, but he was arrested after police chased his car to a wealthy neighborhood. A three-hour gun battle erupted involving more than 100 police and soldiers, and the house was riddled with bullets, but there were no reported deaths. Arellano Felix was taken into custody.	BBC
Arturo Beltrán Leyva	Yes	About 1 week prior to his death, U.S. intelligence revealed information on his location to the Mexican Navy. Killed in a two-hour gun battle during a raid. At least three other cartel operatives were killed during the raid, with a fourth death by suicide. 5 days prior to the raid, there was a capture attempt in the same municipality in which 4 deaths occured, but he escaped. With U.S. intelligence, he was traced to another safe house and captured. 3 gunmen killed along with an innocent bystander	The Guardian Wikileaks
Ignacio Coronel Villareal	Yes	Killed in a short raid in which Villareal began firing at military personnel, killing one soldier. He was then shot dead. There is no further publicly available information on what lead to the death of this kingpin.	CNN The Guardian
Antonio Cárdenas Guillén	No	Mexican armed forces had been trying to capture him for nearly 6 months prior to his death, nearly capturing him on 6 occasions. Because the failed captures could create anticipation effects, I exclude this capture from my analysis.	El Universal
Flavio Mendez Santiago	Yes	There are no publicly available details about what lead to his eventual arrest.	InSight Crime
Jorge Eduardo Costilla Sánchez	No	There are no publicly available details on the intelligence used to locate Sánchez. Because his capture occurred following its onset and occurred in a state where Operation Northern Lynx occurred and occurred following its onset and occurred during Caledrón's presidency, this capture is excluded from my analysis. (see Section 2.3 for a discussion of Operation Northern Lynx).	
Iván Velázquez Caballero	No	Located using intelligence from the DEA. Because his capture occurred in a state where Operation Northern Lynx occurred following its onset and occurred during Caledrón's presidency, this capture is excluded from my analysis. (see Section 2.3 for a discussion of Operation Northern Lynx).	InSight Crime El Universal Lessing (2017)
Miguel Treviño Morales	Yes	A month before his capture, American authorities passed along information that he was making frequent trips near the U.S. border, due to the birth of a newborn son. According to a U.S. official who was not authorized to speak, the authorities traded intelligence obtained via wiretaps and informant tips that led Mexican authorities to Mr. Trevino's truck. Mexican marines in a helicopter intercepted Treviño and arrested him and two aides without a shot.	The New York Times
Mario Ramirez Treviño	Yes	5 days earlier, lower-level members of the cartel were captured. Through this, Mexican authorities learned of Treviño's location and in a joint operation, authorities arrested him without a single shot being fired.	Univision Noticias

Continued on next page

Kingpin Name	Included	Capture/Death Details	Source(s)
Nazario Moreno González	No	The previous administration announced in 2010 that he had been killed in a shootout, but on February 7, 2014, a lower-ranked cartel member that was arrested revealed that he was still alive. On March 4, special forces found high-caliber weapons buried in Tumbiscatío, which led them to a vehicle that Nazario used. They located Nazario shortly after and encircled him and ordered him to surrender, but he refused and started shooting, after which he was killed. Because he was believed to be dead prior to his actual death, it is likely that the rumored death had effects on market structure. Because of this, I exclude this kingpin from my analysis.	Milenio
Héctor Beltrán Leyva	Yes	Captured while dining at a restaurant in a joint operation by the Attorney General's Office and the Ministry of National Defense after 11 months of intelligence work. The arrest was carried out quickly and without shots fired.	Cronica
Vicente Carillo-Fuentes	Yes	Following a surveillance operation, Carrillo was detained at a checkpoint, where he identified himself with a driver's license with a false name. But when he was identified by the authorities, he surrendered without resistance.	Reuters
Servando Gómez Martínez	Yes	After being tipped by locals, who had formed a militia to oppose the cartel's presence, authorities learned of a messenger between Gómez and lower-level cartel members. He was subsequently wiretapped, which narrowed the search down to 10 properties. After his girlfriend delivered chocolate cake to the cave in which he was hiding on February 7, authorities learned of his location and spent a few weeks confirming his location and planning the arrest. He was arrested with no shots fired.	The Tico Times
Abigael González Valencia	Yes	Arrested by the Mexican Navy using intelligence obtained through investigations by the Attorney General with no shots fired. He was arrested promptly once intelligence was obtained (within 24 hours of learning about his whereabouts).	Excelsior 24 Horas
Omar Treviño Morales	Yes	Killed in a pre-dawn raid in an operation with no casualties. Authorities learned of his location through intelligence obtained by tracking a lower-level financial operator.	Animal Politico Milenio
Joaquín Archivaldo Guzmán Loera	No	Guzmán has been captured on 5 occasions, and escaped from Mexico's top security prison on 2 occasions. It is unclear which of these captures to use, and whether events in which he escaped would have effects on market structure. Thus, he is excluded from the analysis.	SOFREP Time WIRED

Notes: This table reports details on the events leading up to the capture of each kingpin identified in the U.S. Kingpin Act Sanctions List during the event window (2008-2015) in chronological order and whether or not they are included in the final sample. Because the 2016 capture is potentially endogenous, the event window is restricted to 2008-2015.

Appendix Table A3. Largest 10 ω Weights for Number of Criminal Enterprises by Treatment Date

Rank	Municipality	Weight	Rank	Municipality	Weight
Cohort: 2008			**Cohort: 2009**		
1	Batopilas, Chia	6.85	1	Cuauhtémoc, CDMX	2.19
2	Bocoyna, Chia	5.90	2	Ecatepec de Morelos, Edomex.	1.16
3	Guazapares, Chia	5.47	3	Mérida, Yuc.	0.97
4	Urique, Chia	4.93	4	Ocoyoacac, Edomex.	0.97
5	Guachochi, Chia	4.83	5	Ixtapan de la Sal, Edomex.	0.79
6	Soyopa, Son.	3.31	6	Mazatlán, Sin	0.79
7	Hualahuises, N.L.	3.11	7	Valle de Bravo, Edomex.	0.79
8	San Javier, Son.	2.98	8	Arcelia, Gro.	0.60
9	Onavas, Son.	2.92	9	Tejupilco, Edomex.	0.60
10	Chihuahua, Chia	2.78	10	Saltillo, Coah	0.60
Cohort: 2010			**Cohort: 2011**		
1	Cuauhtémoc, CDMX	1.07	1	Ocoyoacac, Edomex.	0.38
2	Mérida, Yuc.	0.83	2	Mérida, Yuc.	0.08
3	Ocoyoacac, Edomex.	0.82	3	Colima, Col.	0.07
4	Batopilas, Chia	0.67	4	Ixtapan de la Sal, Edomex.	0.05
5	Ecatepec de Morelos, Edomex.	0.61	5	Miguel Hidalgo, CDMX	0.05
6	Ixtapan de la Sal, Edomex.	0.60	6	Arcelia, Gro.	0.05
7	Valle de Bravo, Edomex.	0.60	7	Tejupilco, Edomex.	0.05
8	Mazatlán, Sin	0.59	8	Valle de Bravo, Edomex.	0.05
9	Arcelia, Gro.	0.53	9	Saltillo, Coah	0.05
10	Tejupilco, Edomex.	0.53	10	Jesús María, Ags.	0.05
Cohort: 2013			**Cohort: 2014**		
1	Veracruz, Ver.	2.14	1	Taxco de Alarcón, Gro.	0.22
2	Toluca, Edomex.	1.11	2	Los Cabos, B.C.S	0.19
3	Chapala, Jal	0.91	3	Manzanillo, Col.	0.19
4	Ayotlán, Jal	0.75	4	Juárez, Chia	0.18
5	Allende, N.L.	0.75	5	Chapala, Jal	0.16
6	Metepec, Edomex.	0.75	6	Tláhuac, CDMX	0.14
7	Tepic, Nay.	0.68	7	Villa Victoria, Edomex.	0.14
8	Armería, Col.	0.67	8	Cuitláhuac, Ver.	0.14
9	Benito Juárez, CDMX	0.65	9	Venustiano Carranza, CDMX	0.14
10	Nicolás Romero, Edomex.	0.64	10	Metepec, Hgo.	0.14
Cohort: 2015.0					
1	Yecapixtla, Mor.	1.45			
2	Petatlán, Gro.	1.39			
3	Hermosillo, Son.	1.20			
4	Delicias, Chia	1.17			
5	Tuxpan, Mich.	1.16			
6	Tixtla de Guerrero, Gro.	0.89			
7	Amacuzac, Mor.	0.88			
8	Piedras Negras, Coah	0.87			
9	Técpan de Galeana, Gro.	0.82			
10	Solidaridad, Q.R.	0.79			

Notes: This table reports the percentage share of the synthetic control unit for the 10 municipalities with the largest ω weights for each treatment date. The weights in this table are for the outcome: number of criminal enterprises.

Appendix Table A4. Largest 10 ω Weights for Number of Already Existing Non-Splinter CEs by Treatment Date

Rank	Municipality	Weight	Rank	Municipality	Weight
Cohort: 2008			**Cohort: 2009**		
1	Batopilas, Chia	20.88	1	San Lucas Quiaviní, Oax.	1.95
2	Bocoyna, Chia	7.34	2	Los Cabos, B.C.S	1.59
3	Chihuahua, Chia	7.29	3	Asientos, Ags.	0.86
4	Guazapares, Chia	6.32	4	Aguascalientes, Ags.	0.04
5	Urique, Chia	5.18	5	Calvillo, Ags.	0.04
6	Guachochi, Chia	4.97	6	Cosío, Ags.	0.04
7	Hualahuises, N.L.	2.93	7	Jesús María, Ags.	0.04
8	San Javier, Son.	2.74	8	Pabellón de Arteaga, Ags.	0.04
9	Soyopa, Son.	2.64	9	Rincón de Romos, Ags.	0.04
10	Miguel Hidalgo, CDMX	2.42	10	San José de Gracia, Ags.	0.04
Cohort: 2010			**Cohort: 2011**		
1	Cuauhtémoc, CDMX	1.42	1	Los Cabos, B.C.S	1.72
2	Juárez, Chia	1.21	2	San Lucas Quiaviní, Oax.	1.23
3	Iztapalapa, CDMX	1.20	3	Asientos, Ags.	0.62
4	Tepic, Nay.	1.03	4	Aguascalientes, Ags.	0.04
5	Naucalpan de Juárez, Edomex.	1.02	5	Calvillo, Ags.	0.04
6	Puebla, Pue.	1.00	6	Cosío, Ags.	0.04
7	Almoloya de Juárez, Edomex.	0.81	7	Jesús María, Ags.	0.04
8	Aguascalientes, Ags.	0.79	8	Pabellón de Arteaga, Ags.	0.04
9	Ecatepec de Morelos, Edomex.	0.79	9	Rincón de Romos, Ags.	0.04
10	Atizapán de Zaragoza, Edomex.	0.62	10	San José de Gracia, Ags.	0.04
Cohort: 2013			**Cohort: 2014**		
1	Lázaro Cárdenas, Tlax.	1.78	1	Taxco de Alarcón, Gro.	5.45
2	Iztapalapa, CDMX	1.14	2	Cuauhtémoc, CDMX	5.13
3	Aguascalientes, Ags.	0.32	3	Asientos, Ags.	0.25
4	Asientos, Ags.	0.05	4	Aguascalientes, Ags.	0.04
5	Calvillo, Ags.	0.05	5	Calvillo, Ags.	0.04
6	Cosío, Ags.	0.05	6	Cosío, Ags.	0.04
7	Jesús María, Ags.	0.05	7	Jesús María, Ags.	0.04
8	Pabellón de Arteaga, Ags.	0.05	8	Pabellón de Arteaga, Ags.	0.04
9	Rincón de Romos, Ags.	0.05	9	Rincón de Romos, Ags.	0.04
10	San José de Gracia, Ags.	0.05	10	San José de Gracia, Ags.	0.04
Cohort: 2015.0					
1	Cuauhtémoc, CDMX	1.43			
2	Chilpancingo de los Bravo, Gro.	1.37			
3	Chihuahua, Chia	1.10			
4	Acapulco de Juárez, Gro.	1.03			
5	Toluca, Edomex.	0.98			
6	Yecapixtla, Mor.	0.95			
7	Mazatlán, Sin	0.91			
8	Delicias, Chia	0.84			
9	Bocoyna, Chia	0.74			
10	Tuxpan, Mich.	0.71			

Notes: This table reports the percentage share of the synthetic control unit for the 10 municipalities with the largest ω weights for each treatment date, along with the names of the associated municipalities. The weights in this table are for the outcome: number of non-new splinter criminal enterprises that were present in the year before treatment in municipality m.

Appendix Table A5. Largest 10 ω Weights for Number of New Non-Splinter CEs by Treatment Date

Rank	Municipality	Weight	Rank	Municipality	Weight
Cohort: 2008			**Cohort: 2009**		
1	Batopilas, Chia	2.13	1	Los Cabos, B.C.S	0.45
2	Bocoyna, Chia	1.87	2	San Lucas Quiaviní, Oax.	0.34
3	Guazapares, Chia	1.74	3	Asientos, Ags.	0.17
4	Urique, Chia	1.66	4	Aguascalientes, Ags.	0.05
5	Guachochi, Chia	1.65	5	Calvillo, Ags.	0.05
6	Cuauhtémoc, CDMX	1.32	6	Cosío, Ags.	0.05
7	Soyopa, Son.	1.12	7	Jesús María, Ags.	0.05
8	Moris, Chia	1.11	8	Pabellón de Arteaga, Ags.	0.05
9	Maguarichi, Chia	1.10	9	Rincón de Romos, Ags.	0.05
10	Salinas Victoria, N.L.	1.05	10	San José de Gracia, Ags.	0.05
Cohort: 2010			**Cohort: 2011**		
1	Delicias, Chia	1.80	1	Los Cabos, B.C.S	0.29
2	Tanhuato, Mich.	0.99	2	San Lucas Quiaviní, Oax.	0.25
3	Los Cabos, B.C.S	0.97	3	Mier y Noriega, N.L.	0.21
4	San Lucas Quiaviní, Oax.	0.93	4	San Antonio Nanahuatípam, Oax.	0.16
5	Chihuahua, Chia	0.89	5	Santa Catarina, Gto.	0.13
6	Acapulco de Juárez, Gro.	0.84	6	Delicias, Chia	0.11
7	Nezahualcóyotl, Edomex.	0.83	7	Asientos, Ags.	0.10
8	Atotonilco de Tula, Hgo.	0.80	8	Matehuala, S.L.P.	0.09
9	Miahuatlán de Porfirio Díaz, Oax.	0.78	9	Trancoso, Zac.	0.08
10	Mier y Noriega, N.L.	0.76	10	San Francisco de Conchos, Chia	0.07
Cohort: 2013			**Cohort: 2014**		
1	Acapulco de Juárez, Gro.	0.39	1	Metepec, Edomex.	1.66
2	Batopilas, Chia	0.34	2	Juárez, Chia	1.61
3	Cuauhtémoc, CDMX	0.17	3	Cuauhtémoc, CDMX	1.51
4	Salvatierra, Gto.	0.10	4	Zacatecas, Zac.	1.08
5	Metepec, Edomex.	0.10	5	Chimalhuacán, Edomex.	1.03
6	Juárez, Chia	0.10	6	Benito Juárez, Q.R.	1.01
7	Quecholac, Pue.	0.10	7	Cuautitlán Izcalli, Edomex.	1.01
8	Chihuahua, Chia	0.10	8	Chapala, Jal	0.95
9	Nezahualcóyotl, Edomex.	0.10	9	Ayotlán, Jal	0.94
10	Benito Juárez, CDMX	0.10	10	Acapulco de Juárez, Gro.	0.93
Cohort: 2015.0					
1	Naucalpan de Juárez, Edomex.	2.44			
2	Cuauhtémoc, CDMX	2.15			
3	Tultitlán, Edomex.	2.14			
4	Iztapalapa, CDMX	2.14			
5	Saltillo, Coah	2.11			
6	Aguascalientes, Ags.	2.11			
7	Huixquilucan, Edomex.	1.84			
8	Miguel Hidalgo, CDMX	1.83			
9	Manzanillo, Col.	1.54			
10	Atenco, Edomex.	1.52			

Notes: This table reports the percentage share of the synthetic control unit for the 10 municipalities with the largest ω weights for each treatment date, along with the names of the associated municipalities. The weights in this table are for the outcome: number of non-new splinter criminal enterprises that were not present in the year before treatment in municipality m.

Appendix Table A6. Largest 10 ω Weights for New Splinters of CEs Present in Year Before Treatment by Treatment Date

Rank	Municipality	Weight	Rank	Municipality	Weight
Cohort: 2008			**Cohort: 2009**		
1	Hermosillo, Son.	27.39	1	Cananea, Son.	0.76
2	Juárez, Chia	8.00	2	Nogales, Son.	0.53
3	Aguascalientes, Ags.	0.04	3	Acapulco de Juárez, Gro.	0.14
4	Asientos, Ags.	0.03	4	Aguascalientes, Ags.	0.05
5	Calvillo, Ags.	0.03	5	Asientos, Ags.	0.05
6	Cosío, Ags.	0.03	6	Calvillo, Ags.	0.05
7	Jesús María, Ags.	0.03	7	Cosío, Ags.	0.05
8	Pabellón de Arteaga, Ags.	0.03	8	Jesús María, Ags.	0.05
9	Rincón de Romos, Ags.	0.03	9	Pabellón de Arteaga, Ags.	0.05
10	San José de Gracia, Ags.	0.03	10	Rincón de Romos, Ags.	0.05
Cohort: 2010			**Cohort: 2011**		
1	Cuauhtémoc, CDMX	6.59	1	Campeche, Camp.	0.41
2	Acapulco de Juárez, Gro.	4.51	2	Paraíso, Tab.	0.26
3	Chihuahua, Chia	3.72	3	Veracruz, Ver.	0.08
4	Juárez, Chia	2.52	4	Victoria, Tamps.	0.05
5	Mazatlán, Sin	2.17	5	Saltillo, Coah	0.05
6	Veracruz, Ver.	2.08	6	Ojocaliente, Zac.	0.05
7	Gustavo A. Madero, CDMX	2.07	7	San Antonio Nanahuatípam, Oax.	0.05
8	Aguascalientes, Ags.	1.69	8	Doctor Arroyo, N.L.	0.05
9	Petatlán, Gro.	1.68	9	Atarjea, Gto.	0.05
10	Chilpancingo de los Bravo, Gro.	1.67	10	San Francisco de Conchos, Chia	0.05
Cohort: 2013			**Cohort: 2014**		
1	Veracruz, Ver.	2.81	1	Temoaya, Edomex.	3.23
2	Toluca, Edomex.	1.77	2	Benito Juárez, CDMX	0.79
3	Miguel Hidalgo, CDMX	1.36	3	Juárez, Chia	0.61
4	Chalco, Edomex.	1.09	4	Aguascalientes, Ags.	0.04
5	Atizapán de Zaragoza, Edomex.	1.07	5	Asientos, Ags.	0.04
6	Boca del Río, Ver.	1.05	6	Calvillo, Ags.	0.04
7	Huixquilucan, Edomex.	0.95	7	Cosío, Ags.	0.04
8	Ixtapaluca, Edomex.	0.84	8	Jesús María, Ags.	0.04
9	Cuauhtémoc, CDMX	0.82	9	Pabellón de Arteaga, Ags.	0.04
10	Allende, N.L.	0.72	10	Rincón de Romos, Ags.	0.04
Cohort: 2015.0					
1	Yecapixtla, Mor.	3.00			
2	Tuxpan, Mich.	1.60			
3	Amacuzac, Mor.	1.54			
4	Tlatlaya, Edomex.	1.53			
5	Solidaridad, Q.R.	1.52			
6	Amatepec, Edomex.	1.51			
7	Tixtla de Guerrero, Gro.	1.50			
8	Zihuatanejo de Azueta, Gro.	0.82			
9	Batopilas, Chia	0.81			
10	Miguel Hidalgo, CDMX	0.80			

Notes: This table reports the percentage share of the synthetic control unit for the 10 municipalities with the largest ω weights for each treatment date, along with the names of the associated municipalities. The weights in this table are for the outcome: number of new splinter criminal enterprises of CEs that were present in the year before treatment in municipality m.

Appendix Table A7. Largest 10 ω Weights for New Splinters of CEs Not Present in Year Before Treat by Treatment Date

Rank	Municipality	Weight	Rank	Municipality	Weight
Cohort: 2008			**Cohort: 2009**		
1	La Piedad, Mich.	16.21	1	Cuauhtémoc, CDMX	8.70
2	Gustavo A. Madero, CDMX	4.99	2	Petatlán, Gro.	4.91
3	Xochimilco, CDMX	0.12	3	Juan R. Escudero, Gro.	4.54
4	Aguascalientes, Ags.	0.04	4	Zihuatanejo de Azueta, Gro.	4.44
5	Asientos, Ags.	0.04	5	Villa Aldama, Ver.	4.31
6	Calvillo, Ags.	0.04	6	Chilpancingo de los Bravo, Gro.	4.29
7	Cosío, Ags.	0.04	7	Ameca, Jal	4.18
8	Jesús María, Ags.	0.04	8	Tlatlaya, Edomex.	4.14
9	Pabellón de Arteaga, Ags.	0.04	9	Atenango del Río, Gro.	4.11
10	Rincón de Romos, Ags.	0.04	10	Azcapotzalco, CDMX	4.07
Cohort: 2010			**Cohort: 2011**		
1	La Piedad, Mich.	1.01	1	Santa María Huatulco, Oax.	2.52
2	Juan R. Escudero, Gro.	0.58	2	Hermosillo, Son.	0.59
3	Toluca, Edomex.	0.42	3	Ocoyoacac, Edomex.	0.35
4	Aguascalientes, Ags.	0.05	4	Aguascalientes, Ags.	0.04
5	Asientos, Ags.	0.05	5	Asientos, Ags.	0.04
6	Calvillo, Ags.	0.05	6	Calvillo, Ags.	0.04
7	Cosío, Ags.	0.05	7	Cosío, Ags.	0.04
8	Jesús María, Ags.	0.05	8	Jesús María, Ags.	0.04
9	Pabellón de Arteaga, Ags.	0.05	9	Pabellón de Arteaga, Ags.	0.04
10	Rincón de Romos, Ags.	0.05	10	Rincón de Romos, Ags.	0.04
Cohort: 2013			**Cohort: 2014**		
1	Chapala, Jal	1.22	1	Chapala, Jal	0.58
2	Casas Grandes, Chia	0.71	2	Venustiano Carranza, CDMX	0.40
3	Venustiano Carranza, CDMX	0.70	3	Casas Grandes, Chia	0.40
4	Tecámac, Edomex.	0.56	4	Metepec, Edomex.	0.37
5	Nicolás Romero, Edomex.	0.52	5	Tecomán, Col.	0.36
6	Las Choapas, Ver.	0.44	6	Tepic, Nay.	0.35
7	Nanchital de Lázaro Cárdenas del Río, Ver.	0.43	7	Nicolás Romero, Edomex.	0.34
8	Ojocaliente, Zac.	0.42	8	Batopilas, Chia	0.30
9	Salvatierra, Gto.	0.41	9	Santa María del Oro, Jal	0.28
10	Calimaya, Edomex.	0.41	10	Santa María del Oro, Nay.	0.27
Cohort: 2015.0					
1	Ecatepec de Morelos, Edomex.	0.35			
2	Batopilas, Chia	0.19			
3	Chapala, Jal	0.11			
4	Bocoyna, Chia	0.10			
5	Guazapares, Chia	0.09			
6	Manzanillo, Col.	0.09			
7	Urique, Chia	0.09			
8	Guachochi, Chia	0.09			
9	Casas Grandes, Chia	0.08			
10	Nicolás Romero, Edomex.	0.08			

Notes: This table reports the percentage share of the synthetic control unit for the 10 municipalities with the largest ω weights for each treatment date, along with the names of the associated municipalities. The weights in this table are for the outcome: number of new splinter criminal enterprises of CEs that were not present in the year before treatment in municipality m.

Appendix Table A8. Largest 10 ω Weights for Homicides Per 100,000 Population by Treatment Date

Rank	Municipality	Weight	Rank	Municipality	Weight
Cohort: 2008-Q1			**Cohort: 2008-Q4**		
1	Batopilas, Chia	0.20	1	Santiago Lalopa, Oax.	3.40
2	Bocoyna, Chia	0.13	2	Santa María Cortijo, Oax.	0.87
3	Guachochi, Chia	0.12	3	Santa María Tataltepec, Oax.	0.71
4	Guazapares, Chia	0.08	4	Aguascalientes, Ags.	0.04
5	Urique, Chia	0.08	5	Asientos, Ags.	0.04
6	Yécora, Son.	0.08	6	Calvillo, Ags.	0.04
7	Santiago Lalopa, Oax.	0.07	7	Cosío, Ags.	0.04
8	Moris, Chia	0.07	8	Jesús María, Ags.	0.04
9	San Juan Juquila Mixes, Oax.	0.07	9	Pabellón de Arteaga, Ags.	0.04
10	Ajuchitlán del Progreso, Gro.	0.07	10	Rincón de Romos, Ags.	0.04
Cohort: 2009-Q4			**Cohort: 2010-Q3**		
1	Praxedis G. Guerrero, Chia	0.54	1	Hidalgo, Dgo.	1.57
2	San Mateo Nejápam, Oax.	0.52	2	General Treviño, N.L.	0.45
3	Santa María Tataltepec, Oax.	0.22	3	Onavas, Son.	0.22
4	Aguascalientes, Ags.	0.05	4	Aguascalientes, Ags.	0.05
5	Asientos, Ags.	0.05	5	Asientos, Ags.	0.05
6	Calvillo, Ags.	0.05	6	Calvillo, Ags.	0.05
7	Cosío, Ags.	0.05	7	Cosío, Ags.	0.05
8	Jesús María, Ags.	0.05	8	Jesús María, Ags.	0.05
9	Pabellón de Arteaga, Ags.	0.05	9	Pabellón de Arteaga, Ags.	0.05
10	Rincón de Romos, Ags.	0.05	10	Rincón de Romos, Ags.	0.05
Cohort: 2011-Q1			**Cohort: 2013-Q3**		
1	Batopilas, Chia	1.48	1	Oquitoa, Son.	3.27
2	Jiménez, Tamps.	1.24	2	Sáric, Son.	1.79
3	Praxedis G. Guerrero, Chia	1.16	3	San Pedro Molinos, Oax.	1.51
4	Aguascalientes, Ags.	0.04	4	Aguascalientes, Ags.	0.04
5	Asientos, Ags.	0.04	5	Asientos, Ags.	0.04
6	Calvillo, Ags.	0.04	6	Calvillo, Ags.	0.04
7	Cosío, Ags.	0.04	7	Cosío, Ags.	0.04
8	Jesús María, Ags.	0.04	8	Jesús María, Ags.	0.04
9	Pabellón de Arteaga, Ags.	0.04	9	Pabellón de Arteaga, Ags.	0.04
10	Rincón de Romos, Ags.	0.04	10	Rincón de Romos, Ags.	0.04
Cohort: 2014-Q4			**Cohort: 2015-Q1**		
1	Los Cabos, B.C.S	31.84	1	San Juan Chicomezúchil, Oax.	0.97
2	San Juan Chicomezúchil, Oax.	2.76	2	Santa María del Rosario, Oax.	0.96
3	Maguarichi, Chia	0.36	3	San Juan Ñumí, Oax.	0.78
4	Aguascalientes, Ags.	0.03	4	Aguascalientes, Ags.	0.05
5	Asientos, Ags.	0.03	5	Asientos, Ags.	0.05
6	Calvillo, Ags.	0.03	6	Calvillo, Ags.	0.05
7	Cosío, Ags.	0.03	7	Cosío, Ags.	0.05
8	Jesús María, Ags.	0.03	8	Jesús María, Ags.	0.05
9	Pabellón de Arteaga, Ags.	0.03	9	Pabellón de Arteaga, Ags.	0.05
10	Rincón de Romos, Ags.	0.03	10	Rincón de Romos, Ags.	0.05

Notes: This table reports the percentage share of the synthetic control unit for the 10 municipalities with the largest ω weights for each treatment date, along with the names of the associated municipalities. The weights in this table are for the outcome: number of homicides per 100,000 population

Appendix Table A9. Largest 10 ω Weights for Homicides Per 100,000 Population for Males Aged 15-39 by Treatment Date

Rank	Municipality	Weight	Rank	Municipality	Weight
Cohort: 2008-Q1			**Cohort: 2008-Q4**		
1	Batopilas, Chia	34.46	1	Cucurpe, Son.	2.48
2	Bocoyna, Chia	21.48	2	Santiago Laxopa, Oax.	2.43
3	Guazapares, Chia	17.21	3	San Mateo Nejápam, Oax.	0.92
4	Guachochi, Chia	15.42	4	Aguascalientes, Ags.	0.04
5	Urique, Chia	8.10	5	Asientos, Ags.	0.04
6	Maguarichi, Chia	3.12	6	Calvillo, Ags.	0.04
7	Arizpe, Son.	0.20	7	Cosío, Ags.	0.04
8	Aguascalientes, Ags.	0.00	8	Jesús María, Ags.	0.04
9	Asientos, Ags.	0.00	9	Pabellón de Arteaga, Ags.	0.04
10	Calvillo, Ags.	0.00	10	Rincón de Romos, Ags.	0.04
Cohort: 2009-Q4			**Cohort: 2010-Q3**		
1	Praxedis G. Guerrero, Chia	0.49	1	San Bernardo, Dgo.	0.73
2	San Mateo Nejápam, Oax.	0.25	2	Praxedis G. Guerrero, Chia	0.30
3	Aguascalientes, Ags.	0.05	3	Onavas, Son.	0.27
4	Asientos, Ags.	0.05	4	Doctor Coss, N.L.	0.15
5	Calvillo, Ags.	0.05	5	Aguascalientes, Ags.	0.05
6	Cosío, Ags.	0.05	6	Asientos, Ags.	0.05
7	Jesús María, Ags.	0.05	7	Calvillo, Ags.	0.05
8	Pabellón de Arteaga, Ags.	0.05	8	Cosío, Ags.	0.05
9	Rincón de Romos, Ags.	0.05	9	Jesús María, Ags.	0.05
10	San José de Gracia, Ags.	0.05	10	Pabellón de Arteaga, Ags.	0.05
Cohort: 2011-Q1			**Cohort: 2013-Q3**		
1	Praxedis G. Guerrero, Chia	2.42	1	Tubutama, Son.	0.48
2	Maguarichi, Chia	1.19	2	Mina, N.L.	0.35
3	Guadalupe, Chia	0.63	3	Santiago del Río, Oax.	0.33
4	Aguascalientes, Ags.	0.04	4	Otzoloapan, Edomex.	0.29
5	Asientos, Ags.	0.04	5	Tepemaxalco, Pue.	0.25
6	Calvillo, Ags.	0.04	6	Apaxtla, Gro.	0.22
7	Cosío, Ags.	0.04	7	López, Chia	0.21
8	Jesús María, Ags.	0.04	8	Gran Morelos, Chia	0.21
9	Pabellón de Arteaga, Ags.	0.04	9	Santo Domingo Xagacía, Oax.	0.19
10	Rincón de Romos, Ags.	0.04	10	Chinicuila, Mich.	0.19
Cohort: 2014-Q4			**Cohort: 2015-Q1**		
1	Gran Morelos, Chia	10.47	1	Bocoyna, Chia	8.34
2	San Antonio Nanahuatípam, Oax.	7.55	2	Oquitoa, Son.	0.78
3	Maguarichi, Chia	0.36	3	Gran Morelos, Chia	0.27
4	Aguascalientes, Ags.	0.04	4	Aguascalientes, Ags.	0.04
5	Asientos, Ags.	0.04	5	Asientos, Ags.	0.04
6	Calvillo, Ags.	0.04	6	Calvillo, Ags.	0.04
7	Cosío, Ags.	0.04	7	Cosío, Ags.	0.04
8	Jesús María, Ags.	0.04	8	Jesús María, Ags.	0.04
9	Pabellón de Arteaga, Ags.	0.04	9	Pabellón de Arteaga, Ags.	0.04
10	Rincón de Romos, Ags.	0.04	10	Rincón de Romos, Ags.	0.04

Notes: This table reports the percentage share of the synthetic control unit for the 10 municipalities with the largest ω weights for each treatment date, along with the names of the associated municipalities. The weights in this table are for the outcome: number of homicides of males aged 15-39 per 100,000 population

Appendix Table A10. Largest 10 ω Weights for Homicides Per 100,000 Population, Excluding Males Aged 15-39 by Treatment Date

Rank	Municipality	Weight	Rank	Municipality	Weight
Cohort: 2008-Q1			**Cohort: 2008-Q4**		
1	Batopilas, Chia	19.00	1	Santiago Lalopa, Oax.	1.39
2	Bocoyna, Chia	13.74	2	Santa María Tataltepec, Oax.	0.32
3	Guachochi, Chia	11.70	3	Santa María Cortijo, Oax.	0.29
4	Urique, Chia	10.99	4	Aguascalientes, Ags.	0.05
5	San Juan Cieneguilla, Oax.	9.77	5	Asientos, Ags.	0.05
6	Moris, Chia	8.35	6	Calvillo, Ags.	0.05
7	Guazapares, Chia	6.78	7	Cosío, Ags.	0.05
8	Cosoltepec, Oax.	5.50	8	Jesús María, Ags.	0.05
9	San Antonio Nanahuatípam, Oax.	4.89	9	Pabellón de Arteaga, Ags.	0.05
10	Santiago Lalopa, Oax.	3.39	10	Rincón de Romos, Ags.	0.05
Cohort: 2009-Q4			**Cohort: 2010-Q3**		
1	Coahuayutla de José María Izazaga, Gro.	0.95	1	San Nicolás Hidalgo, Oax.	3.38
2	Santiago Nundiche, Oax.	0.84	2	Cusihuiriachi, Chia	0.38
3	Coyame del Sotol, Chia	0.43	3	Onavas, Son.	0.17
4	Aguascalientes, Ags.	0.05	4	Aguascalientes, Ags.	0.04
5	Asientos, Ags.	0.05	5	Asientos, Ags.	0.04
6	Calvillo, Ags.	0.05	6	Calvillo, Ags.	0.04
7	Cosío, Ags.	0.05	7	Cosío, Ags.	0.04
8	Jesús María, Ags.	0.05	8	Jesús María, Ags.	0.04
9	Pabellón de Arteaga, Ags.	0.05	9	Pabellón de Arteaga, Ags.	0.04
10	Rincón de Romos, Ags.	0.05	10	Rincón de Romos, Ags.	0.04
Cohort: 2011-Q1			**Cohort: 2013-Q3**		
1	Atil, Son.	4.93	1	Santiago Zoochila, Oax.	0.71
2	Soyopa, Son.	1.17	2	Oquitoa, Son.	0.56
3	San Juan Lajarcia, Oax.	0.91	3	Asunción Tlacolulita, Oax.	0.34
4	Aguascalientes, Ags.	0.04	4	Bacanora, Son.	0.33
5	Asientos, Ags.	0.04	5	Gran Morelos, Chia	0.32
6	Calvillo, Ags.	0.04	6	Mina, N.L.	0.30
7	Cosío, Ags.	0.04	7	El Oro, Dgo.	0.29
8	Jesús María, Ags.	0.04	8	Rosario, Chia	0.29
9	Pabellón de Arteaga, Ags.	0.04	9	Riva Palacio, Chia	0.28
10	Rincón de Romos, Ags.	0.04	10	Sáric, Son.	0.28
Cohort: 2014-Q4			**Cohort: 2015-Q1**		
1	San Andrés Tepetlapa, Oax.	11.36	1	Güémez, Tamps.	2.13
2	San Juan Chicomezúchil, Oax.	1.57	2	Santa María Yolotepec, Oax.	0.90
3	Trincheras, Son.	0.43	3	San Francisco Huehuetlán, Oax.	0.69
4	Aguascalientes, Ags.	0.04	4	Aguascalientes, Ags.	0.04
5	Asientos, Ags.	0.04	5	Asientos, Ags.	0.04
6	Calvillo, Ags.	0.04	6	Calvillo, Ags.	0.04
7	Cosío, Ags.	0.04	7	Cosío, Ags.	0.04
8	Jesús María, Ags.	0.04	8	Jesús María, Ags.	0.04
9	Pabellón de Arteaga, Ags.	0.04	9	Pabellón de Arteaga, Ags.	0.04
10	Rincón de Romos, Ags.	0.04	10	Rincón de Romos, Ags.	0.04

Notes: This table reports the percentage share of the synthetic control unit for the 10 municipalities with the largest ω weights for each treatment date, along with the names of the associated municipalities. The weights in this table are for the outcome: number of homicides per 100,000 population, excluding males aged 15-39.

Appendix Table A11. Effect of Kingpin Captures on Non-Organized Crime Article Volume

	Pre-Treatment Mean		Level Change		% Change	
	(1)	(2)	(3)	(4)	(5)	(6)
Number of Non-Organized Crime Articles Mentioning Location	2637.09	2578.18	-105.70	-106.06	-4.01	-4.11
	(2186.69)	(2164.92)	(147.47)	(147.47)	(5.59)	(5.72)
Controls	No	Yes	No	Yes	No	Yes
Number of Cohorts	7	7	7	7	7	7
N Treated	12	12	12	12	12	12
N Control	2,151	2,151	2,151	2,151	2,151	2,151

Notes: This table reports estimates of the effect of kingpin captures and deaths on the number of unique newspaper articles that are predicted by the classification mode to not be about organized crime that mention a given municipality in a given year. Columns (1) and (2) show the cohort-weighted and SDID weighted pre-treatment mean, with standard errors reported below. Column (1) uses control variables during the selection of the SDID weights, and column (2) does not. Columns (3) and (4) show the estimated treatment effects in levels with and without controls, respectively. The weighting process for the synthetic control units is described in Appendix C.1. Columns (5) and (6) show the percentage change in the outcome relative to the pre-treatment mean, weighted by the SDID weights, as defined in Equation 8. Controls include year indicators interacted with indicies measuring the suitability of municipality m for cultivating opium poppies and cannabis, respectively, as discussed in 3.4. Standard errors for coefficient estimates are obtained through block bootstrapping with 100 replications as outlined in Clarke et al. (2023). All regressions include year and municipality fixed effects. * p<0.1, ** p<0.05, *** p<0.01.

Appendix Table A12. Correlation of Treatment Status With Corruption Proxy

	Homicide Gap Per 100,000 Population	
	(1)	(2)
Kingpin Captured/Killed	44.53	4.94
	(44.13)	(49.66)
Constant	175.03***	173.01***
	(5.88)	(6.04)
Control for Baseline Criminal Presence	No	Yes
N	2163	2163

Notes: Column 1 reports coefficients from the regression of the municipal homicide gap, aggregated between 2015 and 2019 and standardized by average population in those years, on an indicator variable that is equal to 1 when municipality m had a kingpin capture or death during the event window. The municipal homicide gap is defined as the difference between homicides reported by the INEGI, which are derived from death certificates and homicides reported by SESNSP, which are derived from data reported by state Attorney General's offices and reflect cases that were actually prosecuted by the government. Column 2 reports coefficients from the same regression but controls for baseline criminal presence, defined as the the average number of criminal enterprises in municipality m on an annual basis between 2005-2008, before the event window begins. *p<0.1; **p<0.05; ***p<0.01

Appendix Table A13. Change in Market Concentration (HHI) Between 2005 and 2019

	Mean HHI		Change: 2019 – 2005	
	2005	2019	Level Change	Std. Error
Herfindahl-Hirschman Index (HHI)	0.776	0.702	-0.074 (t = -2.02)	0.037 (p = 0.046)
N (2005) N (2019)	72 259			

Notes: This table reports the difference in average adjusted HHI of criminal enterprises between 2005 and 2019. The reported change reflects the difference in means, with standard errors computed using Welch's unequal variance t-test. The t-statistic and p-value are shown below the estimated change. Observations correspond to municipalities with non-missing HHI values in the respective years. * p<0.1, ** p<0.05, *** p<0.01.

Appendix Table A14. Effect of Kingpin Captures and Deaths on Market Structure (Including Potentially Endogenous Events)

	Pre-Treatment Mean		Level Change		% Change	
	(1)	(2)	(3)	(4)	(5)	(6)
Number of Criminal Enterprises	2.35	2.35	0.91***	0.91***	38.97***	38.58***
	(1.73)	(1.74)	(0.29)	(0.29)	(12.32)	(12.28)
Number of Non-New Splinter CEs Present in Year Before Treatment	1.20	1.22	-0.12	-0.12	-9.85	-9.68
	(1.32)	(1.32)	(0.16)	(0.16)	(12.91)	(12.74)
Number of Non-New Splinter CEs Not Present in Year Before Treatment	0.16	0.16	0.65***	0.64***	399.26***	400.93***
	(0.52)	(0.52)	(0.11)	(0.11)	(66.61)	(67.18)
Number of New Splinters of CEs Present in Year Before Treatment	0.33	0.33	0.40***	0.40***	119.75***	119.28***
	(0.61)	(0.61)	(0.11)	(0.11)	(32.62)	(32.62)
Number of New Splinters of CEs Not Present in Year Before Treatment	0.35	0.35	0.03	0.03	9.01	8.46
	(0.60)	(0.60)	(0.07)	(0.07)	(18.56)	(18.54)
Controls	No	Yes	No	Yes	No	Yes
Number of Cohorts	9	9	9	9	9	9
N Treated	17	17	17	17	17	17
N Control	2,151	2,151	2,151	2,151	2,151	2,151

Notes: The reported coefficients are estimated effects of kingpin captures or deaths on outcomes related to market structure, as estimated using Synthetic Difference-in-Differences and defined in Equation 6. This version of the results includes all kingpin captures and deaths, including potentially endogenous ones that were excluded from the main analysis. The number of criminal enterprises includes CEs of all types (cartels, factions, and independent groups). I define new splinters as either new factions of a criminal enterprise (CE) or former factions that recently split from a CE (within the past 4 years). Columns (1) and (2) show the cohort-weighted and SDID weighted pre-treatment mean, with standard errors reported below. Column (1) uses control variables during the selection of the SDID weights, and column (2) does not. Columns (3) and (4) show the estimated treatment effects in levels with and without controls, respectively. The weighting process for the synthetic control units is described in Appendix C.1. Columns (5) and (6) show the percentage change in the outcome relative to the pre-treatment mean, weighted by the SDID weights, as defined in Equation 8. Controls include year indicators interacted with indicies measuring the suitability of municipality m for cultivating opium poppies and cannabis, respectively, as discussed in 3.4. Standard errors for coefficient estimates are obtained through block bootstrapping with 100 replications as outlined in Clarke et al. (2023). All regressions include year and municipality fixed effects. * p<0.1, ** p<0.05, *** p<0.01.

Appendix Table A15. Effect of Kingpin Captures and Deaths on Homicide Rates (Including Potentially Endogenous Events)

	Pre-Treatment Mean		Level Change		% Change	
	(1)	(2)	(3)	(4)	(5)	(6)
Homicides Per 100,000 Population	5.15	5.24	2.15***	1.58***	41.81***	30.23***
	(4.83)	(4.87)	(0.50)	(0.50)	(9.71)	(9.55)
Homicides Per 100,000 Population (Males Aged 15-39)	2.80	3.08	1.77***	1.36***	63.43***	44.27***
	(2.77)	(2.99)	(0.38)	(0.38)	(13.49)	(12.25)
Homicides Per 100,000 Population (Excluding Males Aged 15-39)	2.42	2.33	0.86***	0.68***	35.65***	29.27***
	(3.43)	(3.49)	(0.26)	(0.26)	(10.81)	(11.24)
Controls	No	Yes	No	Yes	No	Yes
Number of Cohorts	12	12	12	12	12	12
N Treated	17	17	17	17	17	17
N Control	2,151	2,151	2,151	2,151	2,151	2,151

Notes: The reported coefficients are estimated effects of kingpin captures or deaths on outcomes related to homicide rates per 100,000 population, as estimated using Synthetic Difference-in-Differences and defined in Equation 6. This version of the results includes all kingpin captures and deaths, including potentially endogenous ones that were excluded from the main analysis. I report estimates for the overall homicide rates, in addition to decomposing the results according to whether or not the victim is in a demographic group that is likely to be involved in organized crime. This includes males 15-39, as discussed by Calderón et al. (2015). Columns (1) and (2) show the cohort-weighted and SDID weighted pre-treatment mean, with standard errors reported below. Column (1) uses control variables during the selection of the SDID weights, and column (2) does not. Columns (3) and (4) show the estimated treatment effects in levels with and without controls, respectively. The weighting process for the synthetic control units is described in Appendix C.1. Columns (5) and (6) show the percentage change in the outcome relative to the pre-treatment mean, weighted by the SDID weights, as defined in Equation 8. Standard errors for coefficient estimates are obtained through block bootstrapping with 100 replications as outlined in Clarke et al. (2023). All regressions include year-quarter and municipality fixed effects. Additional controls include year-quarter indicators interacted with indicies measuring the suitability of municipality m for cultivating opium poppies and cannabis, respectively, as discussed in 3.4. Standard errors are clustered at the state level to account for correlation in treatment status within states. All regressions include year-quarter and municipality fixed effects. * p<0.1, ** p<0.05, *** p<0.01.

Appendix Table A16. Effect of Kingpin Captures and Deaths on Market Structure (No Imputation)

	Pre-Treatment Mean		Level Change		% Change	
	(1)	(2)	(3)	(4)	(5)	(6)
Number of Criminal Enterprises	1.33	1.35	1.20***	1.19***	90.11***	87.74***
	(1.21)	(1.22)	(0.33)	(0.33)	(24.85)	(24.44)
Number of Non-New Splinter CEs Present in Year Before Treatment	0.34	0.34	0.04	0.04	10.62	10.79
	(0.50)	(0.50)	(0.04)	(0.04)	(12.08)	(12.35)
Number of Non-New Splinter CEs Not Present in Year Before Treatment	0.43	0.40	0.92***	0.92***	215.58***	229.58***
	(0.74)	(0.73)	(0.24)	(0.24)	(55.64)	(59.88)
Number of New Splinters of CEs Present in Year Before Treatment	0.13	0.13	0.17	0.17	136.11	135.21
	(0.37)	(0.37)	(0.12)	(0.12)	(98.81)	(98.41)
Number of New Splinters of CEs Not Present in Year Before Treatment	0.39	0.39	0.22*	0.22*	57.35*	56.96*
	(0.69)	(0.69)	(0.12)	(0.12)	(31.96)	(31.82)
Controls	No	Yes	No	Yes	No	Yes
Number of Cohorts	7	7	7	7	7	7
N Treated	12	12	12	12	12	12
N Control	2,151	2,151	2,151	2,151	2,151	2,151

Notes: The reported coefficients are estimated effects of kingpin captures or deaths on outcomes related to market structure, as estimated using Synthetic Difference-in-Differences and defined in Equation 6. These results use a version of the data in which there is no imputation of criminal enterprise presence between years in a municipality. The number of criminal enterprises includes CEs of all types (cartels, factions, and independent groups). I define new splinters as either new factions of a criminal enterprise (CE) or former factions that recently split from a CE (within the past 4 years). Columns (1) and (2) show the cohort-weighted and SDID weighted pre-treatment mean, with standard errors reported below. Column (1) uses control variables during the selection of the SDID weights, and column (2) does not. Columns (3) and (4) show the estimated treatment effects in levels with and without controls, respectively. The weighting process for the synthetic control units is described in Appendix C.1. Columns (5) and (6) show the percentage change in the outcome relative to the pre-treatment mean, weighted by the SDID weights, as defined in Equation 8. Controls include year indicators interacted with indicies measuring the suitability of municipality m for cultivating opium poppies and cannabis, respectively, as discussed in 3.4. Standard errors for coefficient estimates are obtained through block bootstrapping with 100 replications as outlined in Clarke et al. (2023). All regressions include year and municipality fixed effects. * p<0.1, ** p<0.05, *** p<0.01.

Appendix Table A17. Effect of Kingpin Captures and Deaths on Market Structure (Including Neighboring Jurisdictions in Treated Units)

	Pre-Treatment Mean		Level Change		% Change	
	(1)	(2)	(3)	(4)	(5)	(6)
Number of Criminal Enterprises	4.12	4.18	0.73*	0.73*	17.71*	17.53*
	(2.98)	(3.06)	(0.39)	(0.39)	(9.41)	(9.28)
Number of Non-New Splinter CEs Present in Year Before Treatment	1.77	1.77	-0.38	-0.38	-21.68	-21.57
	(1.79)	(1.79)	(0.23)	(0.23)	(13.21)	(13.17)
Number of Non-New Splinter CEs Not Present in Year Before Treatment	0.96	0.99	1.01***	1.02***	105.03***	102.50***
	(1.33)	(1.49)	(0.33)	(0.33)	(34.72)	(33.65)
Number of New Splinters of CEs Present in Year Before Treatment	0.89	0.89	0.30	0.30	34.00	33.48
	(0.89)	(0.88)	(0.30)	(0.30)	(33.53)	(33.61)
Number of New Splinters of CEs Not Present in Year Before Treatment	0.59	0.58	-0.19	-0.19	-32.92	-33.52
	(0.69)	(0.68)	(0.13)	(0.13)	(22.57)	(22.90)
Controls	No	Yes	No	Yes	No	Yes
Number of Cohorts	7	7	7	7	7	7
N Treated	12	12	12	12	12	12
N Control	2,151	2,151	2,151	2,151	2,151	2,151

Notes: The reported coefficients are estimated effects of kingpin captures or deaths on outcomes related to market structure, as estimated using Synthetic Difference-in-Differences and defined in Equation 6. The number of criminal enterprises includes CEs of all types (cartels, factions, and independent groups). In this version of the results, jurisdictions neighboring treated municipalities are no longer excluded from the control units and instead treated as though they were part of the municipality where the kingpin capture or death occurred. I define new splinters as either new factions of a criminal enterprise (CE) or former factions that recently split from a CE (within the past 4 years). Columns (1) and (2) show the cohort-weighted and SDID weighted pre-treatment mean, with standard errors reported below. Column (1) uses control variables during the selection of the SDID weights, and column (2) does not. Columns (3) and (4) show the estimated treatment effects in levels with and without controls, respectively. The weighting process for the synthetic control units is described in Appendix C.1. Columns (5) and (6) show the percentage change in the outcome relative to the pre-treatment mean, weighted by the SDID weights, as defined in Equation 8. Controls include year indicators interacted with indicies measuring the suitability of municipality m for cultivating opium poppies and cannabis, respectively, as discussed in 3.4. Standard errors for coefficient estimates are obtained through block bootstrapping with 100 replications as outlined in Clarke et al. (2023). All regressions include year and municipality fixed effects. * p<0.1, ** p<0.05, *** p<0.01.

Appendix Table A18. Effect of Kingpin Captures and Deaths on Homicide Rates (Including Neighboring Jurisdictions in Treated Units)

	Pre-Treatment Mean		Level Change		% Change	
	(1)	(2)	(3)	(4)	(5)	(6)
Homicides Per 100,000 Population	6.23	6.24	1.43	0.54	22.99	8.59
	(8.57)	(8.56)	(1.12)	(1.12)	(17.97)	(17.95)
Homicides Per 100,000 Population (Males Aged 15-39)	3.89	4.08	0.95	0.45	24.43	11.02
	(6.11)	(6.12)	(0.70)	(0.70)	(17.93)	(17.07)
Homicides Per 100,000 Population (Excluding Males Aged 15-39)	2.30	2.20	0.66	0.28	28.68	12.52
	(2.72)	(2.77)	(0.47)	(0.47)	(20.62)	(21.50)
Controls	No	Yes	No	Yes	No	Yes
Number of Cohorts	9	9	9	9	9	9
N Treated	12	12	12	12	12	12
N Control	2,151	2,151	2,151	2,151	2,151	2,151

Notes: The reported coefficients are estimated effects of kingpin captures or deaths on outcomes related to homicide rates per 100,000 population, as estimated using Synthetic Difference-in-Differences and defined in Equation 6. In this version of the results, jurisdictions neighboring treated municipalities are no longer excluded from the control units and instead treated as though they were part of the municipality where the kingpin capture or death occurred. I report estimates for the overall homicide rates, in addition to decomposing the results according to whether or not the victim is in a demographic group that is likely to be involved in organized crime. This includes males 15-39, as discussed by Calderón et al. (2015). Columns (1) and (2) show the cohort-weighted and SDID weighted pre-treatment mean, with standard errors reported below. Column (1) uses control variables during the selection of the SDID weights, and column (2) does not. Columns (3) and (4) show the estimated treatment effects in levels with and without controls, respectively. The weighting process for the synthetic control units is described in Appendix C.1. Columns (5) and (6) show the percentage change in the outcome relative to the pre-treatment mean, weighted by the SDID weights, as defined in Equation 8. Standard errors for coefficient estimates are obtained through block bootstrapping with 100 replications as outlined in Clarke et al. (2023). All regressions include year-quarter and municipality fixed effects. Additional controls include year-quarter indicators interacted with indicies measuring the suitability of municipality m for cultivating opium poppies and cannabis, respectively, as discussed in 3.4. Standard errors are clustered at the state level to account for correlation in treatment status within states. All regressions include year-quarter and municipality fixed effects. * p<0.1, ** p<0.05, *** p<0.01.

B Machine Learning

B.1 Text Classification Model

To begin, I randomly sample 5021 articles from Spanish language articles in the ProQuest Latin American collection, which contains over 21 million documents. For each article, I read the translated version and determine whether or not it is about organized crime, and then apply a label to the article. I split the articles into an 80% training set and a 20% testing set and train a text classification algorithm that produces a binary prediction about whether or not the article is about crime. For the purposes of this task, I label any article that mentions illicit activity or criminal enterprise as a positive instance.

Prior to training the text classifier, it was necessary to correct for class imbalance, which occurred because the positive class (articles about organized crime) are a small proportion of all newspaper articles. In my random sample, the ratio of negative to positive class examples was 13.47 to 1. Without balancing the classes, the model can be highly accurate by simply predicting that all articles are of the negative class, thus failing to detect all articles about organized crime. However, in this case, the article would suffer from a low precision (how often the model is correct when predicting the target class)[69] and recall (the percentage of true organized crime articles found by the model)[70]. To balance the classes, I use a common approach called synonym replacement (Rizos et al., 2019). I create duplicates of the positive class until the classes are balanced but replace 1 word per sentence with a synonym in the duplicated articles. This is done to prevent overfitting my model to particular examples in my training set to ensure it can still accurately classify unseen examples.[71] I also use the F1 score (a function of both recall and precision) to evaluate my model's performance when training the model as accuracy, while important, is not sufficient for a performant model in this context.

Because machine learning algorithms cannot work with raw text directly, the text must be transformed into a matrix composed of vectors representing individual words, called a word embedding. A word embedding captures the relationships between words. For example, in an English-language embedding, the vector representation of the word 'car' would have a smaller euclidean distance to the vector for 'truck' than the vector for 'computer'.

To obtain these embeddings, I use a pre-trained word embedding trained on Spanish-language newspaper articles to split each article into individual words (called tokens in NLP). Let $i = 1...n$ index the set of all unique words across all articles. For each token x_i, I embed sub-word features such as the word's suffix into an embedding table E_F that is indexed using a hash function and uses Bloom embeddings, which are described in Miranda et al. (2022).[72] I use the tokens as input to a neural network with the following layers:

[69] Precision $= \frac{TP}{TP+FP}$

[70] Recall $= \frac{TP}{TP+FN}$

[71] Prior to training the model, the classes are perfectly balanced. However, a small number of documents (49/9348) could not be processed by SpaCy and are thus excluded from the analysis.

[72] I use the default features proposed by the SpaCy package: the norm, prefix, suffix, shape, and an indicator of whether the token has a trailing space. The embeddings for these features have 5000, 2500, 2500, 2500 and 100 rows, respectively. Each vector in the embedding table is of dimension 96.

B.1.1 1. Embedding Layer

For each unique token x_i, I do the following:

- **Embeddings for Sub-word Features**:
  For each feature F, I obtain the embedding for the attribute:

  $$\boldsymbol{h}_F(x_i) = E_F * Hash(x_i, F)$$

  where $Hash(w_i, F)$ maps the feature value to an embedding index.

- **Pre-Trained Word Vectors**:
  I also obtain the pre-trained word embedding for the token:

  $$\boldsymbol{v}_i^{'} = W_v \cdot \boldsymbol{v}_i$$

  where $\boldsymbol{v}_i \in \mathbb{R}^d$, and $d = 300$ for this particular embedding. $\boldsymbol{W}_v$ is a linear projection matrix that transforms the word vector to match the dimension of the feature embeddings, with weights and biases that are updated during training.

- **Concatenation**:
  Following this, for each article $a \in A$ and all tokens x_{ai}, I concatenate the associated embeddings from the previous 2 steps into a into a single vector $\boldsymbol{e}_{ai} \in \mathbb{R}^{576}$

B.1.2 2. Maxout Layer

The embedding $\boldsymbol{e}_{ai}$ is fed through a maxout activation functions with 3 pieces, each with their own set of weights and biases that are trained by the model to produce a vector $\boldsymbol{m}_{ai} \in \mathbb{R}^{96}$ for each unique token. This introduces non-linearities in the model, which helps it learn complex relationships and patterns in the data.

B.1.3 3. Contextual Encoding Layer

Then, for each token in each article (indexed by a), I obtain the vector representations for the tokens immediately before and after the token and concatenated into a matrix $\boldsymbol{w}_{ai} = [\boldsymbol{m}_{a,i-1}, \boldsymbol{m}_{ai}, \boldsymbol{m}_{a,i+1}]$. A series of convolutions, composed of 4 maxout functions with residual connections, each with 3 pieces and trainable weights, is applied to $\boldsymbol{w}_{ai}$ to obtain a new vector representation $\boldsymbol{r}_{ai} \in \mathbb{R}^{96}$ of the word in context.

B.1.4 4. Text Classification Layer

I obtain a simple vector representation of each article (each with m tokens) through summation following what is known in the NLP literature as a Bag of Words approach:

$$BoW(a) = \sum_{i=1}^{m} \boldsymbol{r}_{ai}$$

Vectors for every article a are combined using a linear combination of weights and biases and used as inputs to a logistic activation function to predict the probability p that each article is about organized crime. The prediction is positive if $p \geq 0.5$.

All weights in the model are updated through backpropagation to minimize the a logistic loss function whose inputs are the labels and predictions in the training set. The learned parameters are applied to the testing data and used to evaluate the model's accuracy. The model performs well, with an F1 score of 98.42%. All of the model's metrics (along with the accuracy metrics for all other machine learning models) are reported in Appendix Table C1 and the confusion matrix for the text classification model is reported in Appendix Figure C2. In my dataset, there are 738,435 articles about organized crime that were detected by my model.

My model performs well in comparison to similar models that classify topics in newspaper articles. The following figure compares the F1 score of my model to those of Garcia et al. (2024), Saigal and Khanna (2020), Ramraj et al. (2020), and Abbas et al. (2010) and shows that my model's performance exceeds that of these models.[73]

Appendix Figure C1. F1 Scores of News Topic Classification Models

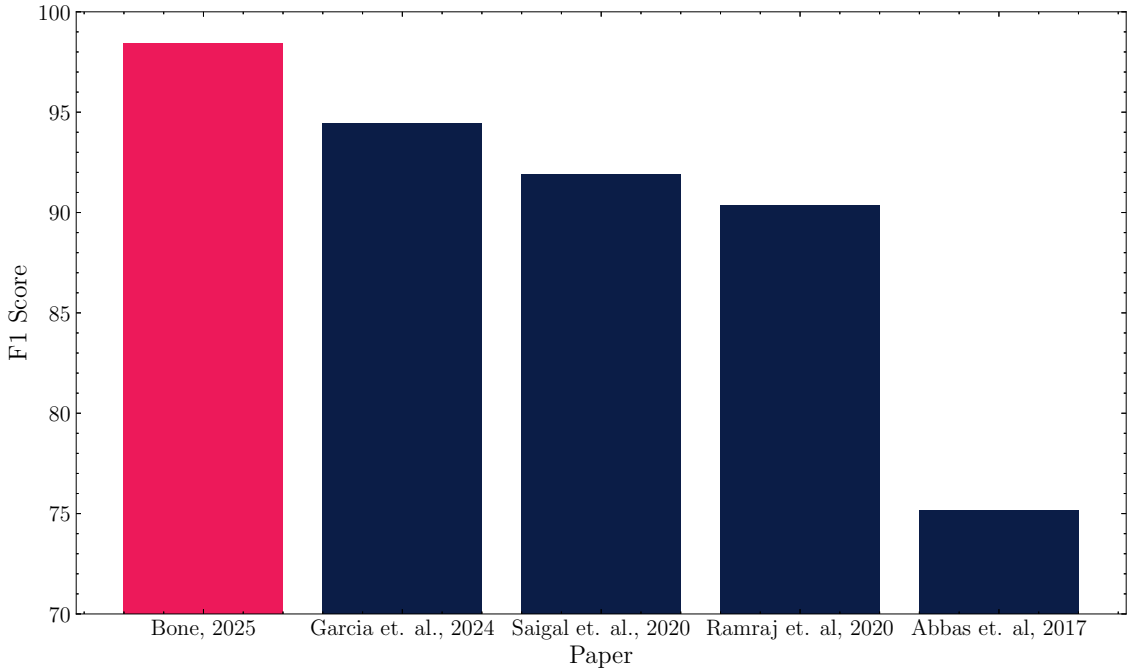

Notes: This figure compares the F1 scores of the text classifier used to similar models that predict topic labels for newspaper articles.

[73] Text classifiers work better when a set of vocabulary is always tied to a particular type of article, which is often true of crime articles. This explains why my model outperforms that of other models, which tend to predict less specific categories such as politics.

B.2 Named Entity Recognition Model

To determine which groups operate in which places, I need to be able to identify criminal enterprises and locations in text. To do this, I employ a named entity recognition (NER) model, a type of natural language processing model used to locate named entities in unstructured text. A big advantage of using named entity recognition is that it is able to identify criminal enterprises in text even if the model hasn't seen the specific organization in training. The use of NER is the feature that differentiates my model the most from previous work that uses news articles to build panel data on cartel presence. Given that hundreds of criminal enterprises operate in Mexico, finding groups based on a list of organizations, as is done in previous research, would lead us to miss a lot of criminal groups (as there are no complete lists of criminal enterprises in Mexico). The same argument applies to locations, as my approach is able to capture sub-municipality level locations and translate them back to the municipality level, rather than querying municipalities directly.

Natural language processing uses contextual clues to identify named entities. In my context, criminal enterprises often appear alongside phrases such as 'cell of', 'operates in' and descriptions of violence, and a named entity recognition model uses these types of clues to identify named entities. The same approach is taken for locations. The main idea is that words that appear in similar contexts are likely related in some way.

To train the NER model, I randomly sample 1001 articles from the 738,435 articles about organized crime and label the names of criminal enterprises and locations in the text. I use the United Nations Office on Drugs and Crime (2018) definition of organized crime when labeling: "Organized crime is a continuing criminal enterprise that rationally works to profit from illicit activities that are often in great public demand". I exclude prison gangs, armed resistance groups without involvement in criminal enterprises, and individual traffickers from my definition of organized crime groups.

The first three layers of the NER model are similar to those in the text classification model, discussed in Appendix Section B.1) in that both use a convolutional neural network with residual connections and a maxout activation function (with the same hyperparameters) to generate the embeddings. However, the NER model uses a pre-trained model with a larger set of vocabulary to generate the word embeddings.[74] The models differ substantially in their final layers. I used the NER framework developed by Honnibal (2017) and employ a transition-based parser to process each article's text and generate a prediction of the entity labels for each token (either a location, criminal enterprise, or non-entity) using a stack and a buffer. A stack is a data structure following the last in, first out principle: The last item added to the stack is the first one to be removed from it. A buffer stores the article's tokens until they are ready to be processed. The process can be conceptualized as a sequence of actions (called transitions) that build the entity structure incrementally.

I follow steps 1-3 in Appendix Section B.1 to generate word embeddings $\boldsymbol{w}_{ai}$ that incorporate information from the surrounding words using the same hyperparameters as the text classification model. Then, for each article in the training set $a \in A$, I initiate an empty stack and a buffer containing an ordered list of tokens x_{ai} in the article. Afterwards, a series

[74] This particular pre-trained model has a vocabulary of 500,000 words.

of actions is applied to each x_{ai}. There are three transition actions:

1. **SHIFT**: Move the next $\boldsymbol{w}_{ai}$ from the buffer to the stack
2. **REDUCE**: Assign a label to the token at the top of the stack and remove it from the stack
3. **ENTITY ACTIONS**: If the token is an entity, label it as either a location or criminal enterprise

The parser iteratively performs the following steps: First, the next token in the article is moved from the buffer to the stack. Then, the following are concatenated into an attribute vector $\boldsymbol{f}_{ai}$:

1. The context-independent vector for the token $\boldsymbol{e}_{ai}$
2. The token's embedding that incorporates information from the surrounding words $\boldsymbol{w}_{ai}$
3. The part of speech tag (e.g. noun, adverb, adjective, etc.) for the token, which is predicted using the pre-trained model
4. Dependency embedding $\boldsymbol{d}_{ai}$, predicted using the pre-trained model, which represents the syntactic dependencies of x_{ai}.[75]
5. A vector of integer representations of the steps taken in previous parsing steps, which provides historical context

The attribute vector $\boldsymbol{f}_{ai}$ is then used as input in a feed-forward neural network with a softmax activation function to predict the the next transition. The model's predictions are compared with the labels and the weights are updated through backpropagation to minimize the cross-entropy loss function. The models are evaluated on the testing data and the resulting accuracy metrics are summarized in Appendix Table C1 and the confusion matrix is reported in Appendix Figure C3. The model obtains a F1 score of 89.83% for locations and a score of 88.79% for criminal enterprises. These scores are comparable to those other Spanish-language NER models. For example, the largest Spanish-language NER model in the SpaCy library, which is one of widely used Python libraries for NER, obtains an average F1 score of 0.89 across all entity types (SpaCy, n.d.).

The following are examples of sentences with at least 1 predicted CE and location (which are bolded). Sentences meeting these criteria are used as input to the relation extraction model:

1. "These and other Municipalities of **Tamaulipas** are fought against the **CDG** by the **Zetas**, who have a greater presence in the center of the State, and the **Northeast Cartel (CDN)**, made up of former **Zetas**."

[75] These dependencies are based on the field of linguistics known as dependency grammar, in which words are connected to each other through directed links. The verb is the central unit of each clause, and all other words are connected either directly or indirectly to the verb through directed links called dependencies.

2. "...The FGR learned that there were some Army weapons in the hands of organized crime when On September 7, 2010, he dismantled a cell of **La Resistencia** in **Guadalajara** and seized weapons and explosives that belonged to the militia."

3. "After his arrest, yesterday federal authorities reinforced security in **Jalisco** and **Morelos** in the face of possible violent reactions from the **CJNG**."

4. "...Colombian civil society considered today that the international community should press for a bilateral truce to shield the peace talks held by the Colombian government and the **FARC** guerrilla in **Havana**, **Cuba**."

5. "Ramírez Treviño was arrested at the beginning of this administration, on August 13, 2013, by elements of the National Defense Secretariat in a spectacular deployment by land and air, in **Reynosa**, **Tamaulipas**, a stronghold of the **Gulf Cartel**, who was He attributes the violence in the entity to the control of drug, arms and money transfer routes, in dispute with the **Los Zetas Cartel**."

B.3 Relation Extraction Model

After training the NER model, I applied it to all articles in my dataset that were about organized crime and extracted all 168,111 articles containing at least 1 criminal enterprise and 1 location. However, a cartel and location co-occurring in the same article is not sufficient to imply that a cartel operates in that location. For example, some news articles link a trafficker (and their associated criminal enterprise) to their place of birth. To resolve this issue, I train a relation extraction model to predict whether a criminal enterprise operates in a particular location. One weakness of relation extraction models is that they are only good at detecting relationships between entities that appear close together in text. Thus, I restrict my dataset to the 219,368 sentences (across all articles) that contain at least 1 criminal enterprise and 1 location. From these articles, I randomly sample 1000 articles and attach a label to the relationship between the organization and location if the group operates within it. I train a neural network with several layers to complete this task:

B.3.1 1. Embedding Layer

For each sentence $s \in S$ containing at least 1 location and CE, I obtain pre-trained word embeddings for each token in the sentence $x_{si}, i = 1...m$ from a the pre-trained Spanish-Language transformer model developed by Cañete et al. (2020):[76]

$$\boldsymbol{X} = Transformer(x_{s1}, x_{s2}...x_{sm}) = \{\boldsymbol{x}_{s1}, \boldsymbol{x}_{s2}...\boldsymbol{x}_{sm}\}$$

where $\boldsymbol{x}_{si} \in \mathbb{R}^{768}$

The model uses the state-of-the art BERT (Bidirectional Encoder Representations from Transformers) framework developed by Devlin et al. (2018), which is known for providing dramatic improvements in model performance over previous state-of-the-art models. Unlike most embedding models, BERT reads text bi-directionally, which means BERT processes a word by looking at the context to the left and right of the word simultaneously. Unlike the SpaCy embedding framework used in previous steps, BERT is more comprehensive and considers all of the words in the sentence when embedding context. This makes BERT especially adept at predicting long-range dependencies between words, which is particularly important in the relation extraction context.

The transformer model used here uses a variant of BERT using whole word masking, in which some tokens during training and the model is trained to predict these masked tokens based on their context. This particular model is a deep neural network model with 12 transformer layers, each of which utilizes the attention mechanism developed by Vaswani et al. (2017) to generate a score that indicates how much attention it should give to the other words in the sentence. Words with higher attention scores contribute more to the final representation of the current word. The embeddings are updated during training to fine-tune the model based on the relation extraction task.

[76] This particular pre-trained model is trained on a large Spanish language corpus that utilizes textual data from a variety of sources, including but not limited to Spanish language Wikipedia pages and news commentary.

B.3.2 2. Entity Pair Generation Layer

The model then applies the NER model discussed Section B.2 to identify entities within the sentence $\boldsymbol{E} = \{e_{s1}, e_{s2}, ..., e_{sN}\}$, where each entity e_{sj} consists of one or more tokens.

I obtain the set of all entity pairs that are within 1000 tokens of each other:

$$Pairs = \{(e_{si}, e_{sj} \mid |i - j| \leq 1000\}$$

B.3.3 3. Instance Tensor Creation Layer

For each entity pair (e_{si}, e_{sj}) and for all $s \in S$, I do the following:[77]

- **Token Vectors for Entities**:
  I then create vectors for the entities, which is the set of all token vectors for all tokens in the entity (referred to as the entity's span in NLP):

$$\boldsymbol{T}_{e_{si}} = \{\boldsymbol{x}_{sk} \mid k \in span\ of\ e_{si}\}$$
$$\boldsymbol{T}_{e_{sj}} = \{\boldsymbol{x}_{sk} \mid k \in span\ of\ e_{sj}\}$$

- **Pooling**:
  I then aggregate the token vectors across tokens in the span through mean pooling:

$$\boldsymbol{v}_{e_{si}} = mean(\boldsymbol{T}_{e_{si}})$$

$$\boldsymbol{v}_{e_{sj}} = mean(\boldsymbol{T}_{e_{sj}})$$

- **Instance Tensor**:
  I then create the instance tensor that is used as input to the classification layer:

$$\boldsymbol{v}_{(e_{si},e_{sj})} = [\boldsymbol{v}_{e_{si}}, \boldsymbol{v}_{e_{sj}}]$$

  where $\boldsymbol{v}_{(e_{si},e_{sj})} \in \mathbb{R}^{1536}$

B.3.4 4. Classification Layer

Using the instance tensor as input, I apply the following steps:

- **Linear Transformation**:
  I first project the instance tensor into a lower dimension. The weights and biases used here are updated during model training:

$$\boldsymbol{h}_{(e_{si},e_{sj})} = \boldsymbol{W}\boldsymbol{v}_{(e_{si},e_{sj})} + \boldsymbol{b}$$

  where $\boldsymbol{W} \in \mathbb{R}^{1536}$ and $\boldsymbol{b} \in \mathbb{R}$.

[77] Each unique order of entities is a separate entity pair.

- **Logistic Function**:
  I then input $\boldsymbol{h}_{(e_{si},e_{sj})}$ into the element-wise logistic sigmoid function to predict the probability of each relationship type:

$$\boldsymbol{p}_{(e_{si},e_{sj})} = \sigma(\boldsymbol{h}_{(e_{si},e_{sj})})$$

  where $\boldsymbol{p}_{(e_{si},e_{sj})} \in [0,1]$ represents the probability of e_{si} operating in e_{sj}. The relation is predicted as positive if $\boldsymbol{p}_{(e_{si},e_{sj})} \geq 0.5$

I then obtain the true labels for each entity pair and compute the logistic loss for the predictions. I update the model's weights and biases through backpropagation to minimize the MSE loss function.

The accuracy metrics of the resulting model are summarized in Appendix Table C1 and the confusion matrix is reported in Appendix Figure C4. The F1 score is 78.83%. It is lower than the F1 scores for the text classification and NER models because the relation extraction task is considerably more complex.

Because the difficulty of the classification task is largely dependent on the textual context, the most suitable comparison paper is Sobrino (2019) who uses a relation extraction model to identify whether a cartel operates in a given municipality. While she does not report the F1 score or the confusion matrix that could be used to calculate this, she does report an accuracy score of 0.86. By comparison, my model obtains an accuracy score of 0.92.[78] The large improvement in accuracy score is the result of using a BERT transformer model to generate the embedding matrix, as these models are known for offering dramatic improvement over previous state-of-the-art models.[79] Additionally, although my model's purpose differs dramatically from relation extraction models used in other literatures, my model's metrics are comparable to other published relation extraction models Han et al. (2020). The highest performing model in Han et al. (2020) obtains a F1 score of 89.5 for an English language model that identifies the relationship between entities and destinations. This should be interpreted with caution because the task at hand is different and English language natural language processing models often outperform their counterparts in other languages.

The following are examples of sentences where at least 1 CE is predicted to be operating in at least 1 location. The predicted entities are highlighted and bold and the detected relationships are listed after the quote:

1. "...The bosses of the Tláhuac, **Jalisco Nueva Generación** and **Unión de Tepito** cartels have agreed to a "non-aggression pact" in **Mexico City** and surrounding areas, a journalistic investigation revealed by the Millennium newspaper."

   Detected Relationships: Jalisco Nueva Generación → Mexico City; Unión de Tepito → Mexico City

[78] F1 score is a more informative statistic than accuracy because it is possible for a model to obtain a high accuracy score by being highly accurate at predicting relationships of the negative class while doing a poor job of classifying relationships of the positive class.

[79] BERT was introduced by researchers at Google in October 2018. While Sobrino (2019) does not specify the type of pre-trained model she uses to generate her word embeddings, it is unlikely that pre-trained Spanish BERT models were available in time for the first draft of her working paper, published in October 2019.

2. “He also gave details about the nicknames of the leaders of the **La Familia** cells in seven towns in the south of the **State of Mexico** and areas of **Guerrero**, who served under the orders of ”El Mojarro”, who in turn He worked for a subject who was only identified by Lugo as ”El Pon”, based in the Mexican town of Palmar Chico.”

   Detected Relationships: La Familia → State of Mexico; La Familia → Guerrero

3. “‘El Pata de Queso’ falls, allegedly involved in the massacre of 72 in **San Fernando** In **Tamaulipas**, Martiniano de Jesús N, “El Pata de Queso”, identified as the regional leader of ”**Los Zetas Vieja Escuela**” was arrested; According to the investigations, the man detained today would have coordinated the murder of 72 migrants in San Fernando in 2010.”

   Detected Relationship: Los Zetas Vieja Escuela → Tamaulipas; Los Zetas Vieja Escuela → San Fernando

4. “The governor pointed out that for the moment within the Guanajuato Seguro operation, the weakening of several criminal groups that generate violence in **Guanajuato** has been carried out, among them the **Santa Rosa de Lima Cartel**, which has been dismantled, arresting so far 960 people.”

   Detected Relationship: Santa Rosa de Lima Cartel → Guanajuato

5. “...To maintain his infiltration in the municipal police of **Iguala** and **Cocula**, the organization The criminal group **Guerreros Unidos** (GU) delivered envelopes of 4,000 pesos to each element, and even bought weapons from those corporations, which collaborated in the disappearance of the 43 normalistas from Ayotzinapa.”

   Detected Relationships: Guerreros Unidos → Iguala; Guerreros Unidos → Cocula

B.4 Geoparsing and Data Cleaning

To account for the ever-increasing presence of small organized crime groups in Mexico, a CE-municipality relationship in a given time period (the year of publication) is identified using a minimum of one article. While each stage of the machine learning pipeline has its own classification scores, given that small organizations have a larger effect on classification metrics at the municipality level, it is important to perform extensive data cleaning.

In my case, data cleaning involves removing special characters from CE names, dropping observations where the location is outside of Mexico, linking abbreviated forms of CEs to their full organization name, correcting spelling errors, and replacing synonyms in organization names so that they are not treated as distinct entities. This was facilitated by using the results of approximate string matching (fuzzy matching) to identify variants of organization names.

I manually look at up to 10 of the articles associated with each unique organization left in my dataset. If there is no evidence that the organization is a Mexican organized crime group, I drop it from my dataset. The use of manual cleaning is common in data mining. For example, it is common to use ML to assist with drug development by producing candidate chemical formulations, after which researchers manually review the model's output (Neil et al., 2018). While this does not eliminate the possibility of false positives, as location or the CE-municipality relationship can still be misclassified, this greatly improves my model's performance. For evidence, I use the content of the articles themselves, Google searches of the group's name and the content of the resulting hits. During this process, I also document any known relationships between the groups that appear during my research and whether or not the group became independent from any parent (if so, I note the date).

Geoparsing is the process of extracting toponyms (place names) from text and mapping their string representation to a geographic coordinates. This is a challenging task because toponyms are often ambiguous. In Mexico, for example, there is both a state of Puebla and the town of Hildago, Puebla. To resolve the ambiguity, I create a rule-based algorithm that uses Google Maps API to disambiguate locations.

For each sentence that shows a CE operating in a location, I first check if the location has multiple jurisdiction levels (e.g. Tijuana, Baja California) because querying these terms together will lead to more accurate geoparsing than querying them separately. To do this, I check if there is a comma and space after each location in each sentence containing a cartel that operates in a location. This will return sub-jurisdictions and their parent but will also return sentences that feature a list of unrelated locations that a cartel operates in. Next, I check if the next word is in the list of locations with relationships to cartels in that sentence. If it is not, I query the location name (without the jurisdiction it belongs to). Then, I look to the next word and see if it too is another location. If it is, then it is very likely that we are dealing with a large list of locations instead of a specific location and thus each location in the list is queried. If it is not, then I query the first two locations together.

Once I had the list of locations to query, I needed to disambiguate the toponyms. I start by querying each location's name using Google Places Autocomplete API, which when given a string returns a list of potential location matches along with some useful information about the administrative level of the returned locations in addition to parent jurisdictions.

The locations are returned in the order of their perceived relevance. Because my data also contains information on CEs groups throughout Central and South America, this is useful information that I use for toponym resolution. Furthermore, the API also returns a Place ID that I queried using the Google Geocoding API to retrieve the coordinates of the location.

Next, I used data on the administrative level of the location candidates to help me disambiguate locations. If the location was the only location in the sentence, I assumed that the location referred to the first one suggested by the Google Places Autocomplete API. If the location was outside of Mexico, the cartel-location relationship was dropped from my data. Additionally, if the location was the only one in the sentence and the Place ID was that of the country of Mexico, I dropped the data because my analysis occurs at the municipality level.

If the sentence has more than one location, I first checked if Mexico was listed as a parent jurisdiction for any of the responses for the candidate locations. If not, these locations were not queried. When this was not the case, I thus had ambiguous locations which were resolved in the next step of the geoparsing algorithm.

To disambiguate the remaining locations, I exploited the fact that there are multiple locations that co-occur in the same sentence. For this, I assumed that two locations that co-occur in the same sentence are likely to be geographically proximate. For example, given a sentence that mentions 2 places: Toronto and London. London is ambiguous: It could be referring to London, UK or London, Canada. Because Toronto also appears in the same sentence, and Toronto is closer to London, Canada than London, UK, we would assume that the location refers to London, Canada. With this in mind, I find all possible combinations of coordinates (in each combination, there is 1 coordinate pair per location mentioned in-text). For each pair of coordinates in the combination, I calculate the geodesic distance between the points, and then take the sum of these distances in the coordinate combination. I then repeat this for all possible combinations. The combination of coordinates that minimizes the aggregate geodesic distance are chosen as the coordinates.

A notable exception to this rule is when there are more than 10 locations listed in the sentence. In this rare instance, given limited computing resources, I break the locations into equally sized subgroups of no more than 10 locations, and perform the above procedure for these subgroups. Finally, I drop all CE-location relationships where the location's geography more than 5 kilometers away from Mexico (some port locations are just outside of the shapefile boundaries). Using the jurisdictional boundaries, I map each remaining location to the state it is a part of. Additionally, I drop cases where the Google Places API returned a state as the predicted location. I also drop cases where Tierra Caliente (the name of a region with a lot of CE activity) is the predicted location, as it is too large to link to a specific municipality.

I also geoparse data from the articles in my dataset that are not about organized crime.[80]

[80]Unfortunately, using Google Places API for this task would be prohibitively expensive as there is a far greater number of locations to geocode. Instead, I follow the same process but use the GeoNames gazeteer, an open-source alternative, to geocode these locations. GeoNames is less effective than Google's API at resolving informal or fine-grained locations, such as colloquial neighborhood names or locally used site labels. However, since the non-crime data is used only to measure general coverage and the volume of articles is large, missing a small number of such references is unlikely to meaningfully affect results. Moreover, all locations are ultimately linked to the municipality level, which GeoNames handles reliably.

The purpose of this data is twofold. First, I use the data to demonstrate that my stylized facts are driven by real changes in criminal presence across municipalities, rather than changes in reporting volume across municipalities. While my main results plot the raw data to make them easier to interpret, I present adjusted results, which control for reporting volume, which are discussed in Section 3.6. Additionally, the data is used in a robustness check, where I evaluate the effects of capturing or killing the top leaders of criminal enterprises on market structure and homicide rates. One key concern is that the results may be driven by variation in reporting volumes that are correlated with treatment status, so this data is useful as a falsification test. I discuss these results in Section 5.3.

Appendix Figure C2. Confusion Matrix - Text Classification Model

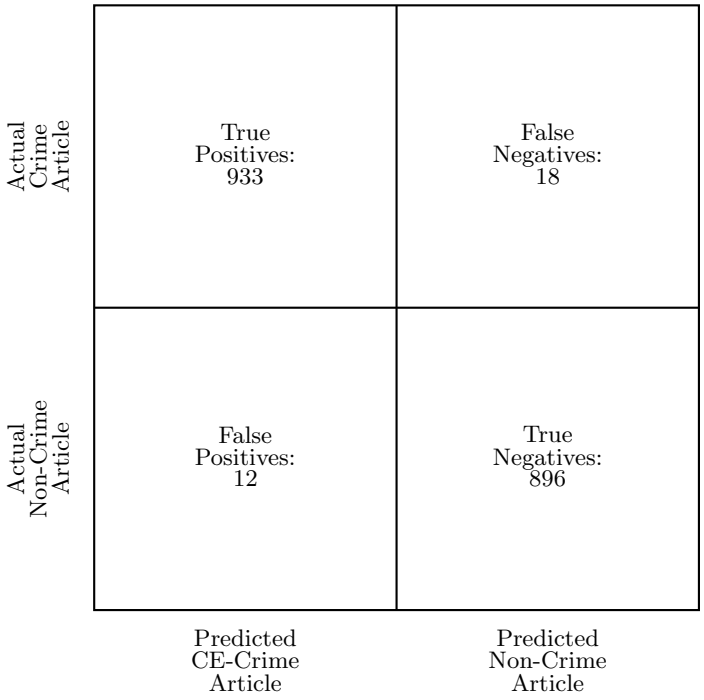

Notes: This figure lists the number of false negatives, true positives, true negatives, and false positives for the testing evaluation data (20% hold out set) used to evaluate the text classification model used to process the data. The model is described briefly in Section 3.1 and in detail in Appendix Section B.1. The model predicts whether a given full-text newspaper article is about organized crime using 5021 hand labeled articles. I label any article with a mention of illicit activity or criminal enterprise as a positive instance. Prior to training, I correct for class imbalance using synonym replacement, as described in Appendix Section B.1.

Appendix Figure C3. Confusion Matrices - Named Entity Recognition

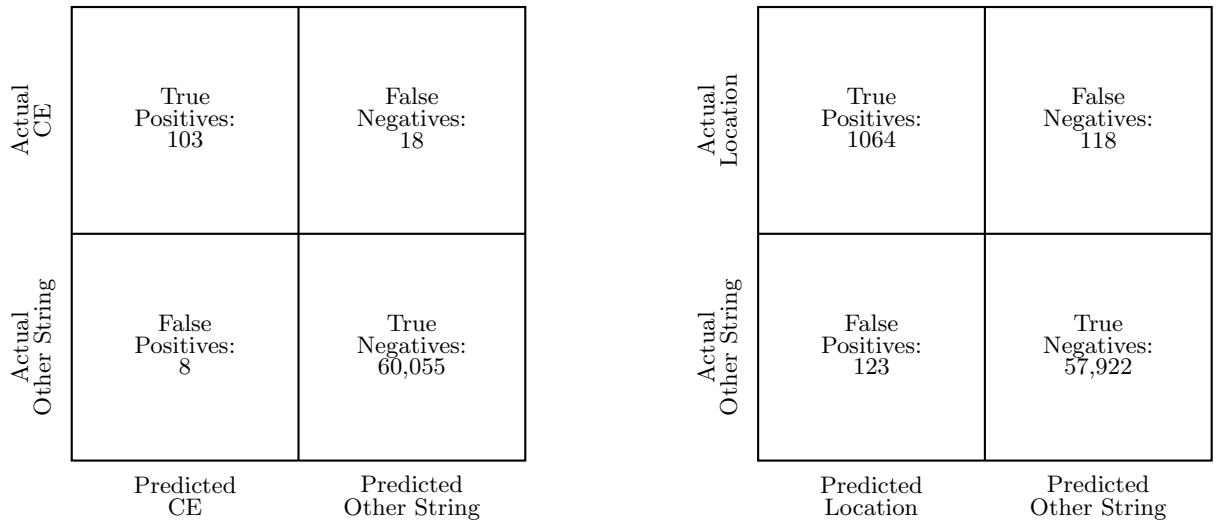

(a) Criminal Enterprises

(b) Locations

Notes: This figure lists the number of false negatives, true positives, true negatives, and false positives for the testing data used to evaluate the Named Entity Recognition model described briefly in Section 3.1 and in detail in Appendix Section B.2. For this model, predictions are made over each possible set of words in articles predicted to be about organized crime, as determined by the text classification model. Because of this, the sum of the entries in the confusion matrix differ from the number of articles in the test set, which is reported in Table C1. Results are presented separately for each of the two entity types (criminal enterprises and locations) according to whether the prediction and actual values were positive or negative. The negative class includes both words (or groups of words) classified as the other type of entity and words that did not have any prediction attached to them.

Appendix Figure C4. Confusion Matrix - Relation Extraction Model

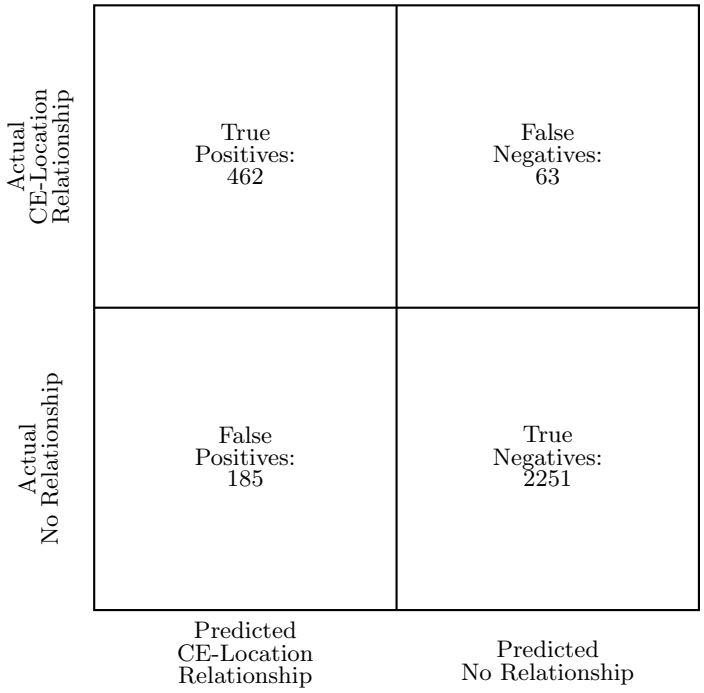

Notes: This figure lists the number of false negatives (FN), true positives (TP), true negatives (TN), and false positives (FP) for the testing data used to evaluate the Relation Extraction model described briefly in Section 3.1 and in detail in Appendix Section B.3. True positives are instances in which the model accurately predicted that a criminal enterprise operated in a given location. True negatives are instances where a criminal enterprise did not operate in a given location and the model did not predict the relationship. False positives are instances where a criminal enterprise was predicted to operate in a given location but did not operate there. False negatives are instances where a criminal enterprise does operate in a given location but the relationship was not predicted by the model. An observation in this case is a potential cartel-CE relationship, which is why the number of observations in the testing data differs from the number of sentences used to test the model's performance, which is reported in Table C1.

Appendix Figure C5. Confusion Matrix - Comparison of CE Panel With Mapping Criminal Organizations Project

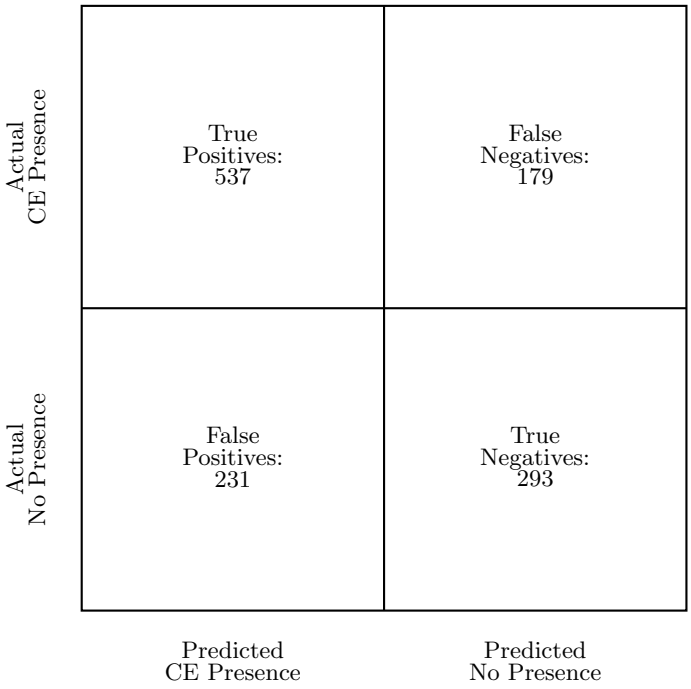

Notes: This figure lists the number of false negatives, true positives, true negatives, and false positives, for the final panel measuring the presence of criminal enterprises using the process described in Section 3.1.1 and detailed in Appendix Section B. I use data from the Mapping Criminal Organizations (MCO) project by Signoret et al. (2021), which is described in Section 3.1.2, as the ground truth for criminal presence. To make the CE panel from this paper comparable with the MCO data, I restrict my panel to cover the same organizations and time period (2007-2015) as the MCO data. Both datasets are aggregated to the state-year level. In the MCO data, a CE is considered to be operating within a state in a given year if it operates in that state in any month in that year and was recorded as having a significant or major presence. This restriction is used to ensure that none of the observations are based on a single source of data, as these observations require agreement between the sources, and some sources use news data. In the CE panel data developed in this paper, a CE is considered to be operating within a given state in a given year if it is operating in at least one municipality within that state. To make my coding consistent with the MCO data, I consider observations listing named factions which existed prior to them ending their affiliation with their parent organization (if applicable) as observations identifying the presence of the larger parent organization. In this context, a prediction is whether a particular CE is operating in a given state in a particular year.

Appendix Table C1. Machine Learning Model Performance by Performance Metric

	Precision	Recall	F1 Score	Observations	Documents (Training)	Documents (Testing)	Matches
Text Classification	98.73	98.11	98.42	21,116,855	7440	1859	738,435
Named Entity Recognition (Locations)	89.64	90.02	89.83	738,435	801	200	168,111
Named Entity Recognition (Criminal Enterprises)	92.79	85.12	88.79	738,435	801	200	168,111
Relation Extraction	71.41	88.00	78.83	219,368	800	200	114,402

Notes: The table above shows prediction accuracy statistics for the main models used in the machine learning pipeline. Additionally, for each model, the table reports the size of the dataset that the model is applied to and how many matches in the dataset the model had. For the text classification model, the model is applied to the entire dataset and the matches refers to the number of articles that were predicted to be about organized crime. For the NER model, the model is applied to all articles that are predicted to be about organized crime, and the matches reflects the number of articles that have at least one criminal enterprise and one location present. The relation extraction model is applied to all sentences that have at least one CE and one criminal enterprises, and the number of matches is the number of sentences where the relation extraction model predicts at least one CE operates in at least one location. The number of training documents is the number of labeled documents used to train each model. In the case of the text classification model, I report the size of the training and testing sets after the synonym replacement, discussed in Appendix Section B.1 is done. For named entity recognition, the number of documents refers to the number of articles that are used to train and test the model. For relation extraction, the number of sentences with both cartels and locations that are used to train the model.

C Econometric Details

C.1 Computation of Synthetic Difference in Difference Weights

As outlined in Arkhangelsky et al. (2021), the vector of cohort-specific unit weights $\hat{\omega}^e$ is part of the solution to the following optimization problem:

$$(\hat{\omega}_0^e, \hat{\omega}_e) = \underset{\omega_0^e \in \mathbb{R}, \omega^e \in \Omega^e}{\arg\min}\ \ell_e(\omega_0^e, \omega^e) \tag{13}$$

where

$$\ell_e(\omega_0^e, \omega^e) = \sum_{t=1}^{T_{pre}} \left(\omega_0^e + \sum_{m=1}^{N_{co}} \omega_m^e Y_{mt} - \frac{1}{N_{tr}^e} \sum_{m=N_{co}+1}^{M} Y_{mt} \right)^2 + \zeta^{2,e} T_{pre} \|\omega^e\|_2^2$$

and

$$\Omega^e = \left\{ \omega^e \in \mathbb{R}_+^N : \sum_{m=1}^{N_{co}} \omega_m^e = 1, \omega_m^e = N_{tr}^{e,-1} \forall m \in N_{co}+1, ..., M \right\}$$

where $\mathbb{R}_+$ denotes the positive real number line and all variables are defined in the same way as in Section 4.2. The regularization parameter ζ^e is given by:

$$\zeta^e = (N_{tr}^e T_{post})^{\frac{1}{4}} \hat{\sigma}^e \tag{14}$$

with

$$\frac{1}{N_{co}(T_{pre}-1)} \sum_{m=1}^{N_{co}} \sum_{t=1}^{T_{pre}-1} (\Delta_{mt} - \bar{\Delta}^e)^2$$

where

$$\Delta_{mt} = Y_{m(t+1)} - Y_{mt}$$

and

$$\bar{\Delta}^e = \frac{1}{N_{co}(T_{pre}+1)} \sum_{m=1}^{N_{co}} \sum_{t=1}^{T_{pre}-1} \Delta_{mt}$$

Unit-specific weights are chosen so that for each cohort, changes in outcomes Y_{mt} in the pre-treatment period are approximately parallel for treatment and control groups. These weights are similar to those proposed by Abadie et al. (2010) for the synthetic control approach but differ in two key ways. First, synthetic DID includes an intercept term ω_0^e in the objective function so that the unit weights $\hat{\omega}_e$ no longer need to make Y_{mt} equal for treatment and control groups in the pre-treatment period: It is enough that their trends be parallel. Second, Arkhangelsky et al. (2021) include a regularization parameter ζ^e to increase the dispersion and ensure the uniqueness of the weights.

Likewise, the cohort-specific vector of time weights, $\hat{\lambda}^e$ is part of the solution to the following optimization problem:

$$(\lambda_0^e, \lambda^e) = \underset{\lambda_0^e \in \mathbb{R}, \lambda^e \in \Lambda^e}{\arg\min}\ \ell_{e,time}(\lambda_0^e, \lambda^e) \tag{15}$$

where

$$\ell_{e,time}(\lambda_0^e, \lambda^e) = \sum_{m=1}^{N_{co}} \left(\lambda_0^e + \sum_{t=1}^{T_{pre}} \lambda_t^e Y_{mt} - \frac{1}{T_{post}} \sum_{t=T_{pre}+1}^{T} Y_{mt} \right)^2$$

$$\Lambda^e = \left\{ \lambda^e \in \mathbb{R}_+^T : \sum_{t=1}^{T_{pre}} \lambda_t^e = 1, \lambda_t^e = T_{post}^{-1} \forall t \in T_{pre} + 1, ..., T \right\}$$

These weights are similar to the unit weights but differ in that they weight only pre-treatment time periods and do not use regularization. Arkhangelsky et al. (2021) show that by including these weights, the estimator can remove bias and include precision by reducing the influence of time periods in which pre-treatment periods for the control group differ greatly from the post-treatment periods. When time-variant controls are included in the specification, the estimation of the weights proceeds as above, except Y_{mt} is replaced with $\tilde{Y}_{mt}$, which is the residual of the regression of Y_{mt} on X_{mt}.