Policing the Hot Hours

High-Frequency Evidence from Hot-Spot Policing in Uruguay

José María Cabrera[a] Alejandro Cid[a] Olivier Marie[b] Federico Veneri[a]

[a]Universidad de Montevideo [b]Erasmus School of Economics and Tinbergen Institute

October 2026

Abstract

Police departments concentrate patrols in high-crime places, but risk also varies sharply within the day. We ask whether allocating patrol jointly across place and hour reduces robbery, how quickly offenders respond when officers arrive and leave, and whether any reduction is offset across nearby hours, locations, or offense types. We study Uruguay's Programa de Alta Dedicación Operativa (PADO), a Montevideo hot-spot program that assigned foot patrols to selected street segments and hours. Combining the operational calendar with GPS measures of realized presence and geocoded crime reports, our benchmark pairs 15,352 patrol episodes in 2016–2018 with the same place, weekday, and clock hour in 2015. During the seven scheduled hours, realized police presence rises by 118 percent and robberies fall by 30 percent relative to their 2015 levels. A placebo that applies the eventual segment-hour schedule to pre-program weeks reproduces neither change, and the decline is not driven by any single start hour. The implied elasticity of robbery with respect to localized patrol exposure is -0.26 to -0.48: a modest response to a very large dose. Robbery responds with a lag: the reduction is about half as large in the first patrol hour as in later hours, and it persists into the hour in which officers return to base. Following the reduction across margins, we detect no rebound in the eight hours after patrol, no rise in the closest neighboring cells at launch, and no shift from robbery to theft; the confidence intervals bound partial displacement rather than exclude it. Patrol allocation has a time dimension, and evaluating place-based policing requires following crime in time as well as space.

JEL codes: K42, H76, R50

Keywords: police, crime, deterrence, hot spots, patrol allocation, high-frequency data, displacement

1 Introduction

Most police officers spend most of their time on patrol, and the central question for a department that cannot hire is where and when to put them. In the economic model of crime, offenders respond to the probability of detection (Becker, 1968), and a visible police officer is the most salient signal of that probability. A large literature confirms that more police means less crime. Police surges after terrorist attacks reduce car theft by block and month in Buenos Aires (Di Tella and Schargrodsky, 2004) and robbery by borough and week in London (Draca et al., 2011); additional resources targeted at robbery reduce it (Machin and Marie, 2011); and changes in force size reduce crime at the city–year level (Chalfin and McCrary, 2018; Mello, 2019). Estimated elasticities range from about -0.1 to -2 depending on the crime and the design (Chalfin, 2025). A second literature knows *where* to send police: hot-spot policing, evaluated in 65 studies, most of them short experiments in US cities, produces small but robust crime reductions in the targeted places, with benefits more often diffusing to adjacent streets than crime displacing to them (Braga et al., 2019).

Much less is known about *when*. If police deter mainly as sentinels rather than as apprehension agents (Nagin et al., 2015), then precisely when officers are visible, and for how long, should matter as much as where they are. Yet the hours at which officers actually stand on a street are rarely observed, and evidence on patrol at high frequency is thin and pulls in opposite directions. In Essex, GPS devices on every officer show that about ten additional minutes of patrol a day in 200-meter circles, for one week, produced a precise zero (Blanes i Vidal and Mastrobuoni, 2018). In Dallas, car GPS shows that when officers are pulled out of their beat by outside calls, a 10 percent fall in presence raises crime by 7 percent (Weisburd, 2021). One study adds very little presence, briefly; the other identifies off losing it. Blanes i Vidal and Mastrobuoni (2018) argue that the gap between these estimates and the large effects of static deployments reflects aggregation: deterrence from a passing patrol decays within about half an hour, so hourly or daily data miss most of it unless presence is sustained.

The question matters beyond the academic puzzle. The US debate about police has centered on staffing: hiring more officers reduces serious crime without raising arrests for serious offenses, but it also produces many more arrests for minor ones, widening the net of the justice system (Chalfin, 2025). Reallocating the hours of the officers a city already has is a different margin. It requires no hiring, and if the mechanism is visibility rather than enforcement, it need not generate arrests at all. Whether that margin delivers, how quickly, and whether the gains are real or merely moved, is what a department needs to know before redrawing its shift pattern. The evidence that could answer these questions is also almost entirely from high-income countries; the review in Chalfin (2025) is restricted to them for lack of comparable work elsewhere, and the largest hot-spot experiment outside them, in Bogotá, found direct effects undone by spillovers (Blattman et al., 2021). Latin America has the highest robbery rates in the world.

This paper studies a program that adds a lot of presence, for hours, on a schedule, and lets us see both the switch-on and the switch-off. The Programa de Alta Dedicación Operativa (PADO), launched in Montevideo in April 2016, assigned pairs of officers on foot to small, high-crime street

segments grouped into administrative circuits. Crucially for our question, the program also assigned patrol hours. At launch the selected segments shared a common late-afternoon start, directing police toward the hours in which robberies were most concentrated; later versions of the program redrew the map ten times and moved scheduled starts across the afternoon and evening. The setting offers unusually detailed measures of both assigned and realized treatment. We reconstruct 15,361 complete scheduled starts between 2016 and 2018, of which 15,352 retain stable segment geography and form the benchmark sample. GPS records from every officer's radio and every vehicle measure realized police presence every five minutes, and geocoded administrative records measure robbery in the same spatial and temporal units.

Our benchmark compares each scheduled patrol episode with the same place, weekday, and clock hour in 2015, and contrasts the seven patrol hours with the four hours just before officers enter. The first difference removes everything permanent about a street and an hour of the week; the second removes anything that changed on that street since 2015 regardless of the hour. What remains is what changes when police step onto the street. Calendar year-month-by-clock-hour effects absorb the ordinary daily cycle of each PADO year relative to 2015.

We ask four questions that correspond to operating decisions a police department has to make. How much does patrol by place and hour reduce robbery, per police-hour? How fast and for how long does robbery respond within the shift, at entry, over seven hours, and at exit? Does it matter at what time the patrol starts? And at what cost elsewhere: does robbery move to later hours, nearby streets, or other offenses?

PADO produced a large and sharply timed change in realized patrol. During the seven scheduled hours, police presence increased by 0.517 hours per segment-hour, about 31 minutes in every hour and 118 percent of its 2015 level, and robberies fell by 1.354 events per 1,000 segment-hours, 30 percent of their 2015 level. The implied elasticity of robbery with respect to localized patrol exposure is -0.26 when both changes are scaled by 2015 levels and -0.48 on a midpoint-arc scale. The effect is large because the dose is large, not because each police-minute is unusually potent. Within the shift, presence is a step function but robbery is not: the reduction in the first patrol hour is about half that of a mid-shift hour, it then plateaus, and it is still present in the transition hour in which presence has fallen to a sixth of its core level. Offenders appear to respond to police with a lag in both directions. The pattern is not an artifact of the clock: scheduled starts range from 14:00 to 18:00, and the core reduction is statistically indistinguishable across start hours. A placebo that applies the eventual schedule to the weeks before PADO existed reproduces neither the patrol increase nor the robbery decline, and the reduction is a similar share of the baseline, 29 to 38 percent, whether or not the highest-risk cells are included.

We then ask whether robbery was deterred or merely displaced. During the eight hours after patrol the estimated rebound is about 10 percent of the core reduction and not distinguishable from zero, although the interval allows nearly half of the reduction to be retimed. Theft does not rise in the treated segments and hours; if anything it falls, and robbery plus theft falls significantly. Around the six synchronized program launches we detect no increase in the closest ring of neighboring

cells, but the spatial comparison is the least precise: the first ring covers many more cells than the treated area, and the estimates do not rule out a meaningful local offset in total robberies. The displacement evidence therefore bounds compensation rather than excludes it.

The paper's main contribution is to make the within-day allocation margin empirically visible in a routine, multi-year program: assigned entry and exit are linked to realized GPS presence and hour-specific robbery, and potential compensation is followed across nearby hours, places, and offense types in the same data. We speak to three literatures. First, the economics of police and crime, which estimates effects of force size or deployment (Di Tella and Schargrodsky, 2004; Draca et al., 2011; Machin and Marie, 2011; Chalfin and McCrary, 2018; Mello, 2019) and, more recently, of routine patrol measured by GPS (Blanes i Vidal and Mastrobuoni, 2018; Weisburd, 2021); PADO adds a large, sustained, scheduled dose of foot patrol under routine conditions, and the within-day margin. Second, work on place-based policing and interference (Braga et al., 2019; Blattman et al., 2021; Ariel et al., 2020; Giulietti et al., 2025), on patrol dosage (Koper, 1995; Williams and Coupe, 2017; Vandeviver et al., 2026) and on predictable deployment (Ariel and Partridge, 2017); PADO adds high-frequency evidence on scheduled entry and exit and on three margins of compensation. Third, earlier evaluations of PADO, which study the program's overall spatial effect at weekly and monthly frequency (Chainey et al., 2021; Cabrera et al., 2022); our object is narrower and more operational, the subdaily allocation of police and what it does hour by hour.

Section 2 describes Montevideo, the mismatch between police and robbery over the day, and PADO. Section 3 presents the data. Section 4 sets out the comparison, the threats to it and the checks we run. Section 5 reports the direct effects and diagnostics. Section 6 asks whether the reduction was deterred or displaced. Section 7 concludes.

2 Setting

2.1 Crime in Montevideo

Uruguay is a small, stable, high-income country for the region, with 3.5 million people. Montevideo, the capital, has 1.4 million, half the country's population and about 80 percent of its robberies (Cabrera et al., 2022). Robbery (*rapiña*) is defined in the penal code as theft by force or threat of force; the use of force is what distinguishes it from theft. Latin America has the highest robbery rates in the world, around 400 per 100,000 inhabitants against about 100 in the United States, and Uruguay's rate rose from under 100 per 100,000 in 1989 to over 900 in 2018.[1] Robbery is also highly concentrated: in 2015, 4.5 percent of the city's 200-by-200-meter grid cells accounted for half of all robberies, a quarter of cells recorded any robbery at all, and the Gini coefficient of robberies across cells was 0.87 (Cabrera et al., 2022). Concentration in time is as marked as concentration in space, with robbery rising from the late afternoon into the night on every day of the week.

Our study period covers the three years after the program's launch. Recorded robberies in

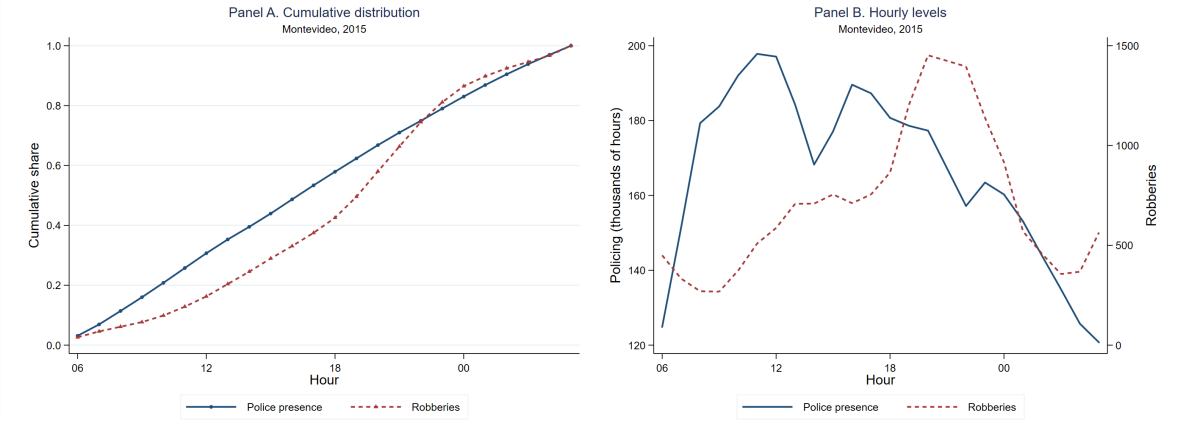
[1] Ministerio del Interior figures as reported in Cabrera et al. (2022). In Montevideo, recorded robberies rose from about 6,000 in 2000 to about 17,000 in 2015, roughly 450 and 1,250 per 100,000 inhabitants.

Montevideo fell from 17,141 in 2015 to 16,166 in 2016 and 15,662 in 2017, then jumped to 23,490 in 2018. The jump follows the entry into force of a new code of criminal procedure on 1 November 2017, which moved Uruguay to an adversarial system and which contemporaneous work associates with a substantial part of the 2018 increase (Cabrera et al., 2022). Our design compares each patrolled street-hour with the same street-hour in 2015 and absorbs calendar month-by-hour differences between each year and 2015, so citywide shifts of this kind are removed before any patrol effect is estimated.

2.2 Police presence and robbery over the day

Figure 1 documents the allocation problem that motivated the program. In 2015, before PADO, police presence measured by GPS was spread almost evenly over the day, whereas robbery rose from the late afternoon into the night. The hours from 06:00 to 18:00 held 53 percent of police time but only 38 percent of robberies. A department that identifies the right streets but keeps the old shift pattern sends officers to the right place at the wrong time.

Figure 1: The daily distribution of police presence and robbery before PADO, Montevideo 2015.



Notes. Panel A plots cumulative shares and Panel B hourly levels. Police presence comes from georeferenced GPS records and robbery from administrative incident records. The figure is descriptive.

2.3 The PADO program

PADO emerged from a broader modernization of the Uruguayan National Police that increased the use of geocoded crime data, crime analysis, and evidence-based deployment. Before PADO, the authorities had experimented with predictive-policing software to identify at-risk areas and with redistributing patrols across groups of street blocks, with mixed results (Cabrera et al., 2022). Under the usual policing model, Montevideo is organized into precincts with territorial jurisdictions, and three features of that model worked against concentrating presence where and when robbery peaks. Foot patrol was one of several tasks in an officer's job description, so demand for other duties pulled officers away from targeted areas. Shifts had fixed start and end times set for administrative

reasons rather than to match the hours of peak crime, and officers had no incentive to apply for the riskier evening window. And previous strategies assigned officers to areas rather than to exact street segments, so that within a high-crime zone officers chose where to stand, and anecdotal evidence suggests they drifted toward the quieter corners of it. PADO was designed to remove each of these constraints: a centralized unit with no precinct boundaries, officers dedicated exclusively to the program, a schedule aimed at the hours of peak robbery, and GPS monitoring of presence on named segments. Many of its officers were drawn from two special units, the Guardia Republicana and the tactical reserve group, whose personnel had not previously been assigned to street patrol.

The program was launched in Montevideo on 11 April 2016 with the objective of reducing robbery by concentrating preventive patrol in the city's highest-risk micro-locations (Inter-American Development Bank and Ministry of Interior of Uruguay, 2018). The initial design contained 120 street segments organized into 28 circuits (Figure 2, left panel). A segment grouped nearby critical points into a patrolable location of roughly 100 to 200 meters. Segments represented less than one percent of the city's urban area but accounted for a disproportionate share of pre-program robbery. Officers patrolled segments on foot, typically in teams of two or three; motorcycles and a supervisor in a patrol vehicle supported the foot patrols at the circuit level. The segment is therefore the substantive treatment location, while the circuit is an administrative and dependence unit.

Initial targeting used geocoded robberies from 2015. A threshold of 15 annual robberies informed the construction of candidate segments, but final assignment also incorporated police knowledge, logistics, and route continuity. Some selected units fell below the threshold because they linked high-crime locations. The assignment rule was thus not a deterministic function of a single running variable, and we do not interpret the threshold as a regression-discontinuity design.

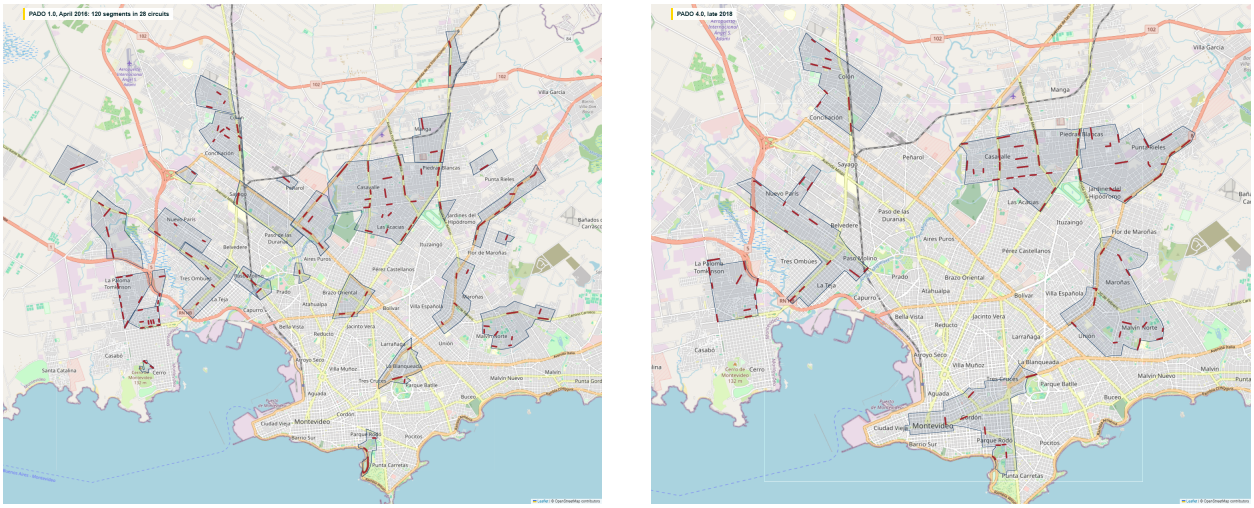
Officers volunteered for PADO and received a monthly bonus of up to 25 percent of base pay in exchange for exclusive dedication to the program and flexible shifts: an officer could work the afternoon shift one week and the night shift the next if the pattern of robbery required it (Cabrera et al., 2022). Compliance was monitored live through the GPS devices in officers' radios. PADO officers were required to leave their segment only to respond to a crime in progress or an emergency call nearby if they were the closest unit, in which case a patrol from the nearest precinct was deployed to maintain presence on the segment; coordination worked well in 2016 and less smoothly in 2017 and 2018. The treatment is therefore existing officers, paid more to work where and when robbery peaks, and we observe both the schedule and the presence actually delivered. The dose was large. During the first version, a treated grid cell received on average about 390 minutes of police presence per day, against about 41 minutes in the rest of the city, and the 257 treated cells covered less than 2 percent of the city's surface but 15 percent of its robberies (Cabrera et al., 2022).

PADO targeted hours as well as places. During the six launch dates, the circuits shared a common 5 p.m. start and operated through the night. Operational information indicates that the final scheduled hour included the return to base rather than a full eighth hour of foot patrol; the realized GPS first stage confirms this distinction. Our event time indexes hours relative to scheduled entry: $k = 0, \ldots, 6$ are the seven core patrol hours and $k = 7$ is a transition hour.

2.4 Program versions and scheduled start times

Both the places and the hours changed over the three years. After the launch, the police's crime analysis unit monitored crime trends continuously and proposed modifications to segments and circuits, and the 2016–2018 administrative files record ten program versions (Figure 2). In July 2016 the program was extended to new segments; in April 2017 the first major redesign removed about 190 of the treated grid cells, added about 160, and kept about 110; the versions that followed first extended coverage and then, from mid-2017 onward, removed segments, so that the program had shrunk by late 2018. On average about 300 cells were patrolled in any version, and 827 of the city's 13,736 cells, 6 percent, were treated at some point (Cabrera et al., 2022). Patrols follow the *relative* hot spots: where robbery falls most, a street drops down the ranking and the next one takes its place. Whether crime also moves is a separate question, which we test directly in Section 6.

Figure 2: PADO foot-patrol segments (red) and circuits (blue): version 1.0, April 2016 (left) and version 4.0, late 2018 (right).



Notes. Operational files, Ministerio del Interior. Base map: OpenStreetMap.

Scheduled entry also varied. Table 1 shows the distribution of scheduled start hours by year. The four most frequent starts are 14:00 to 17:00, but their shares change across years and the 2018 calendar includes a wider set of times. This variation matters for interpretation: relative hour $k = 0$ marks the start of a patrol episode, not a particular hour of the day. It lets us separate hours since police arrived from hour of the day, and ask whether it matters when a patrol starts.

Table 1: Distribution of scheduled start hours

Start hour	2016	2017	2018	Pooled
14:00	1,941 (30.5%)	1,227 (26.3%)	171 (4.0%)	3,339 (21.7%)
15:00	912 (14.3%)	667 (14.3%)	87 (2.0%)	1,666 (10.9%)
16:00	1,981 (31.1%)	783 (16.8%)	1,084 (25.1%)	3,848 (25.1%)
17:00	1,168 (18.3%)	1,167 (25.0%)	806 (18.7%)	3,141 (20.5%)
18:00	0 (0.0%)	233 (5.0%)	858 (19.9%)	1,091 (7.1%)
Other hours	369 (5.8%)	593 (12.7%)	1,305 (30.3%)	2,267 (14.8%)
Total	6,371	4,670	4,311	15,352

Notes. The table uses the stable-geometry benchmark episodes. "Other hours" includes all scheduled starts outside 14:00–18:00. Percentages are within year, except the pooled column. Start times were not randomly assigned.

3 Data

We combine three administrative sources made available by the Ministry of the Interior under confidentiality agreements, each measured at the resolution at which patrol was assigned: the street segment and the hour.

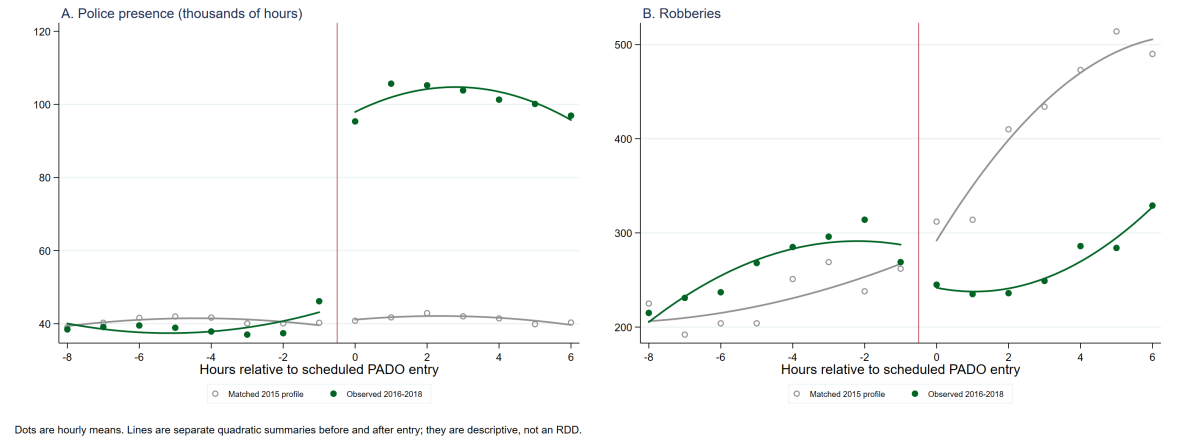
Scheduled patrol. The operational calendar records, for every PADO deployment in 2016–2018, the geographic version, the administrative deployment code, the date, the weekday, and the scheduled start and end hours. From it we reconstruct every scheduled patrol episode, that is, every segment-day on which a foot patrol was assigned, together with its planned entry hour. This is the treatment variable. It is assignment, not delivery, and we keep the two apart throughout.

Realized police presence. Every officer of the Uruguayan National Police carries a TETRA radio with a GPS device, and every vehicle is tracked. The devices transmit their location at high frequency; the Ministry's control room uses the signals to monitor officers and dispatch units. We observe the universe of these signals for Montevideo from 2015 to 2018, roughly 336,000 hours of police presence in an average month, the equivalent of about 1,400 officers on eight-hour shifts every day (Cabrera et al., 2022). We divide time into five-minute bins and count an officer as present in a geographic unit if at least one signal was emitted from inside it during the bin; summing the bins gives police hours per segment-hour, our measure of realized presence. The measure captures officers on foot as well as in vehicles, and it captures them wherever they are, not only where they were sent. Realized presence is an outcome and a first stage rather than a replacement for assignment. This separation avoids conditioning the crime regression on endogenous compliance while showing whether the assigned program changed police activity on the ground. Only a handful of police–crime studies observe presence at all (Sherman and Weisburd, 1995; Blanes i Vidal and Mastrobuoni, 2018; Weisburd, 2021), and none does so at this resolution for a routine, multi-year, city-scale program.

Crime. Geocoded incident records from the police information system, the Sistema de Gestión de Seguridad Pública, provide the location of every reported offense and the time at which the victim states it occurred, which we round down to the hour. We use robbery, theft, and homicide. Robbery is the primary outcome because reducing robbery was PADO's objective and robbery has substantially more hourly support than homicide; theft enters in the offense-substitution exercises. In the matched-2015 core period, robberies account for 39.2 percent of these three recorded offense categories, theft for 60.5 percent, and homicide for 0.3 percent. Because robbery is reported by victims rather than generated by police activity, the outcome is not mechanically tied to the presence of a patrol in the way that arrests or stops would be.

Figure 3 shows what these data look like once aligned on scheduled entry, before any adjustment. Police presence in the treated segments tracks its 2015 profile until the scheduled hour, then jumps; robbery sits below its 2015 profile from the first scheduled hour on. The rest of the paper puts structure and uncertainty on this picture.

Figure 3: Police presence and robbery around scheduled PADO entry, raw profiles.



Notes. The profiles use 15,361 complete scheduled starts. Green dots show observed outcomes in 2016–2018 and gray hollow dots show outcomes in the same geography, weekday, and clock hour in 2015. The curves are separate quadratic summaries before and after scheduled entry; they have no inferential role and the figure imposes no counterfactual adjustment.

Units and samples. We map the operational files into nested but noninterchangeable units. The segment corresponds to the foot-patrol location. A regular 200-by-200-meter cell is the elementary unit used to construct segment outcomes and the spatial spillover graph. Because segment identifiers can be reused across administrative areas and versions, the safe mapping key combines geographic version, deployment code, and segment. The benchmark requires stable segment geography throughout an event. Of 15,361 complete scheduled starts, nine cross-version episodes cannot maintain that mapping and are excluded, leaving 15,352 episodes, 1,353 distinct segment geographies, and 1,503,280 segment-episode-hour observations. The temporal displacement sample requires a clean window from $k = -8$ through $k = 15$, including eight inactive post hours; it contains 13,149 episodes and 63,866 segment-episode pairs. The local spatial design uses the six

launch dates and 2,907 cells in rings zero through three.

The 2015 counterfactual. For each observed PADO date, unit, and clock hour, we construct a pre-program comparison in 2015 using the same ordinal weekday in the calendar. The crosswalk handles week 52/53 explicitly and flags dates without a valid analogue; outcomes are never silently set to zero because of a failed match. Each street-hour is thus its own control, matched to itself one year earlier.

4 Empirical strategy

4.1 The matched-2015 comparison and the phase model

The matched difference is

$$\Delta Y_{sek} = Y_{sek} - Y_{sek}^{2015},$$

where s denotes a segment, e a scheduled episode, and k the hour relative to scheduled entry. The 2015 outcome uses the same segment geography, weekday, and clock hour. The comparison fixes the selected geography and supplies a no-PADO profile for the same place and hour. It does not by itself solve targeting: locations were selected using 2015 crime, and the active hours may have been selected because their 2015 outcomes were unusually high. A street picked because 2015 was bad would tend to look better in 2016 whether or not police arrived.

Our design therefore adds a second difference between the active hours and the hours immediately before officers arrive, on the same street and the same day. If the street simply cooled off relative to 2015, it cooled off at 2 p.m. as well as at 6 p.m. What is left is what changes when police step onto the street. We estimate the phase model

$$\Delta Y_{sek} = \alpha_{se} + \beta_E \, Early_k + \beta_C \, Core_k + \beta_T \, Transition_k + \lambda_{\mathrm{ym}(e,k)\times h(e,k)} + \varepsilon_{sek}.$$

The fixed effect α_{se} absorbs the level of the matched difference within each segment-episode pair, so every comparison is within the same street on the same day. The omitted phase is $k = -4, \ldots, -1$; $Early$ is $k = -8, \ldots, -5$; $Core$ is $k = 0, \ldots, 6$; and $Transition$ is $k = 7$. Calendar year-month-by-clock-hour effects λ absorb aggregate differences between the PADO years and 2015 that vary over the day, so that an ordinary late-night decline cannot mechanically become an effect of elapsed patrol time. The primary coefficient is β_C. Because every pair contributes four baseline and seven core hours, OLS gives equal weight to each segment-episode pair. Police presence is measured in hours per segment-hour and robbery in events per 1,000 segment-hours.

4.2 Identifying assumption and threats

The causal interpretation of β_C requires that, absent PADO, the matched difference between core and baseline hours would not have changed differentially at scheduled entry after accounting for calendar month-by-hour effects. It is a parallel-trends assumption within the day: without patrol,

evening and afternoon on the same street would have changed alike since 2015. It can fail in five ways, and we run a check against each. First, hot streets may cool down after a bad 2015 regardless of patrol; the within-episode comparison with the hours just before arrival, on the same street and the same day, removes any such decline that is common to the day (Section 5.1). Second, the patrol *hours* themselves may have been picked because 2015 evenings were unusually bad on those streets, in which case the decline would be specific to the core hours; a placebo that applies the eventual schedule to the weeks before PADO existed tests this directly (Section 5.5). Third, the pattern could be the evening rather than the police; scheduled starts range from 14:00 to 18:00, so we can compare episodes that reach the same clock hour at different points of the shift (Section 5.4). Fourth, a few extreme streets could drive everything; we drop the highest-risk cells and re-estimate (Section 5.6). Fifth, police or crime may already be shifting before officers arrive; we look four to eight hours earlier and, for the displacement analysis, keep only episodes with eight patrol-free hours on either side (Sections 5.1 and 6.1). Finally, the decline could be real but local: crime may simply have moved to later hours, nearby streets, or other offenses, which Section 6 follows directly. None of these exercises converts scheduled timing into random assignment; they strengthen or delimit the benchmark.

4.3 Inference

We allow observations from the same longitudinal patrol segment to be correlated over time. This is our main inference rule because the segment is the location to which treatment is assigned and it may reappear across many scheduled episodes. The longitudinal key follows a segment across calendar and operational-file changes and yields 206 clusters. Administrative circuits could instead be the relevant source of common shocks, while shocks associated with a particular episode could cut across segments. Appendix Table A2 therefore keeps the coefficient, sample, and fixed effects unchanged and reports three alternative ways to calculate uncertainty: two-way segment–episode clustering, circuit-level clustering, and a leave-one-circuit-out correction designed for the small number of circuit groups. The central conclusions are unchanged under all four rules.

5 Direct effects and diagnostics

5.1 First stage and benchmark estimates

Table 2 reports the controlled phase estimates. During the core, police presence increases by 0.517 hours per segment-hour and robberies fall by 1.354 events per 1,000 segment-hours. These changes are large relative to matched-2015 core means of 0.439 police hours and 4.479 robberies: presence rises by 118 percent, about 31 additional minutes in each hour, and robbery falls by 30 percent. Police presence is also 0.070 hours higher in the early pre-period; the corresponding robbery difference is small and imprecise. The much larger increase at entry shows that the schedule changed realized deployment, while the pre-entry presence difference cautions against describing the first stage as literally beginning from zero. In the transition hour, the police increase has attenuated to

0.084, a fifth of its 2015 level, while the estimated robbery reduction has only roughly halved, to 0.584. We return to this asymmetry below.

Table 2: Benchmark estimates during scheduled patrol

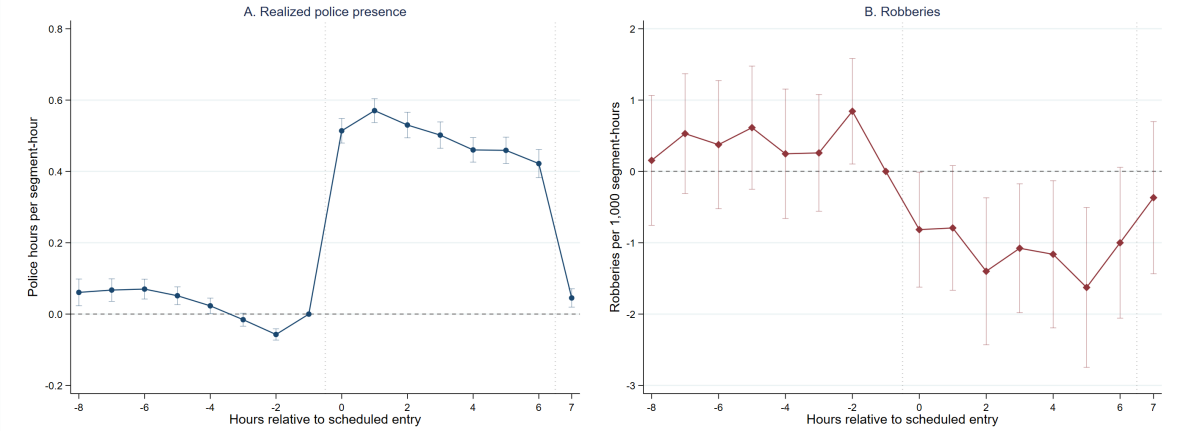
	Police presence		Robberies	
Relative-time phase	Estimate	2015 mean	Estimate	2015 mean
Early pre-period ($k = -8, \ldots, -5$)	0.070***	0.433	0.107	2.190
	(0.009)		(0.232)	
Core patrol hours ($k = 0, \ldots, 6$)	0.517***	0.439	−1.354***	4.479
	(0.018)		(0.281)	
Transition hour ($k = 7$)	0.084***	0.421	−0.584	4.140
	(0.012)		(0.460)	
Observations	1,503,280		1,503,280	
Scheduled episodes	15,352		15,352	
Longitudinal segment clusters	206		206	

Notes. The observation is a segment-by-scheduled-episode-by-relative-hour cell. The dependent variable is the PADO-year outcome minus the outcome in the same segment at the matched ordinal weekday and clock hour in 2015. The omitted phase is $k = -4, \ldots, -1$. Regressions include segment-by-episode and calendar year-month-by-clock-hour fixed effects. Police presence is in police hours per segment-hour and robberies in events per 1,000 segment-hours. Standard errors in parentheses are clustered by longitudinal operational segment. *** $p < 0.01$.

5.2 Within-shift dynamics

Figure 4 traces the hourly coefficients. Presence is a step function: it jumps at scheduled entry, stays roughly flat through the seven core hours, and collapses in the transition hour. Robbery is not. The reduction is smallest in the first patrol hour, deepest in the middle of the shift, and still present in the transition hour once presence has largely gone. Table 3 puts numbers on the figure. Relative to the seven-hour 2015 mean, robbery is about 26 percent lower in hours 0 and 1 and between 30 and 44 percent lower from hour 2 onward; in hour 7 it is 14 percent lower while presence is only a fifth above its 2015 level. The first hour does about half the work of a mid-shift hour, and then the effect plateaus rather than climbs.

Figure 4: Event-study estimates around scheduled PADO entry.



Notes. Points trace hourly coefficients relative to $k = -1$; bars are 95 percent segment-clustered confidence intervals. The vertical dotted lines mark entry and the end of the seven core hours. The sample, outcomes, and fixed effects match Table 2. For robbery, the individual hourly estimates fluctuate before entry rather than forming a sustained pre-trend.

Table 3: The robbery reduction hour by hour

Hours since police arrived	0	1	2	3–4	5	6	7
Robbery, per 1,000 segment-hours	−1.15	−1.13	−1.74	−1.46	−1.96	−1.34	−0.58
Robbery, percent of 2015 level	−26	−25	−39	−33	−44	−30	−14
Police presence, percent of 2015 level	+118 on average, roughly flat over the seven hours						+20
Elasticity by hour (indicative)	−0.13	−0.18	−0.30	−0.20	−0.22	−0.24	≈ −0.7

Notes. Hourly robbery coefficients from the benchmark event study, phase normalization; the 3–4 column averages two hours. Percentages use the seven-hour 2015 core mean (4.479) and, for hour 7, the transition-hour mean (4.140). Police presence is the phase estimate from Table 2.

Two readings are consistent with this pattern. One is that offenders respond to police with a lag in both directions: it takes time for a patrol to be noticed, and deterrence outlasts the patrol once it is. The other is partly mechanical: officers who arrive at twenty past the hour deliver forty minutes of presence in hour 0, and deterrence that decays over half an hour after departure shows up in hour 7 (Blanes i Vidal and Mastrobuoni, 2018). The hourly data cannot separate the two, and we do not claim a gradual build-up of deterrence; the core estimates are not monotonic. We describe the within-shift pattern as a rapid, sustained reduction that starts weaker and ends later than presence itself. Appendix Figure A1 shows the same shape on the clean 24-hour sample, where entry, the seven core hours, the transition, and the eight hours after are all visible in one window.

5.3 Magnitudes and implied elasticities

The first stage establishes that the schedule generated a large change in realized patrol at the intended time. Relative to their matched-2015 means, robberies fall by 30.2 percent while police presence rises by 117.7 percent. The implied elasticity is −0.26 when both changes are scaled by their 2015 means (95 percent interval −0.36 to −0.15) and −0.48 when they are scaled symmetrically around the before and after levels (interval −0.71 to −0.25). The difference between the two is a consequence of the large change in police presence, not a change in sample or model; uncertainty is calculated from a stacked system so that the covariance between the robbery and police-presence estimates enters the ratio.

These statistics scale localized patrol exposure, not citywide staffing. PADO reallocates realized presence across places and hours rather than increasing the size of the police force, and both ratios inherit the identifying assumption of the robbery estimate. With that distinction in mind, Appendix Table A1 places them next to the literature. The PADO elasticity sits at the low end of the range. It is below the city-level robbery elasticity of −0.56 in Chalfin and McCrary (2018) and the −1.3 for violent crime in Mello (2019), both of which scale officers employed rather than officers present. It is below the −0.6 to −0.9 that Weisburd (2021) finds when car patrol is withdrawn from a Dallas beat, and close to the −0.3 to −0.4 from static deployments after terror attacks (Di Tella and Schargrodsky, 2004; Draca et al., 2011), which PADO resembles more than it resembles a roving patrol: officers stand on one segment for seven hours. It is well above the precise zero from ten extra minutes a day in Essex (Blanes i Vidal and Mastrobuoni, 2018). It is also consistent with the earlier evaluation of the same program at weekly frequency, which found an elasticity of −0.47 over three years using cells rather than segments and weeks rather than hours (Cabrera et al., 2022); our midpoint-arc estimate is almost identical, and our baseline-scale estimate is lower because the hourly design measures the dose where and when it was actually delivered. The level effect of PADO is large because presence more than doubles, not because each police-minute is unusually potent. Read together with the Essex result, the comparison suggests diminishing returns to concentration and a role for duration: presence that is too thin or too brief to be noticed does little.

5.4 Elapsed patrol time versus hour of the day

No single frequent start hour drives the pooled pattern. Table 1 shows substantial variation in scheduled entry, and a joint pooled model gives core robbery estimates between −1.07 and −1.43 for starts at 14:00, 15:00, 16:00, and 17:00. The estimates are not statistically distinguishable from one another ($p = 0.961$); Appendix Figure A4 displays them and their uncertainty. Because start times were not randomly assigned and the subgroup intervals are wide, this result does not rule out meaningful heterogeneity or identify the effect of rescheduling patrol from one clock hour to another. It shows instead that the headline pattern is not an artifact of one dominant start time, and that within the afternoon there is no detectably better hour to start.

5.5 Timing placebo

The central selection concern is that the place-hours chosen for PADO might have declined anyway. Hours were chosen on 2015 crime; a high 2015 draw would revert and push ΔY down in the core even without patrol. We ask whether the same timing pattern appears when no PADO patrol was operating. We project the empirical distribution of the benchmark's stable segment geographies, weekdays, and scheduled start times onto compatible dates from 1 January through 9 April 2016. Each pseudo-episode is then compared with the same segment, weekday, and clock hour in 2015. The placebo uses the same seven core hours, four pre-entry hours, weighting logic, fixed effects, and 206 longitudinal segment clusters as the benchmark: every ingredient except patrol.

Before PADO, imposing the eventual schedule produces essentially no change in police presence: the core-versus-pre estimate is -0.004 hours per segment-hour (standard error 0.005). The robbery estimate is -0.211 per 1,000 segment-hours (standard error 0.260; $p = 0.418$). By comparison, the observed PADO estimates are 0.517 for police presence and -1.354 for robbery. The observed-minus-placebo differences are 0.521 hours (standard error 0.020) and -1.143 robberies (standard error 0.385; $p = 0.003$). At the same place-hour resolution as the benchmark, the eventual timing distribution therefore reproduces neither the realized patrol increase nor the robbery decline before PADO operated. The adjustment from -1.354 to -1.143 also shows why the placebo is informative rather than mechanically zero: ordinary patrol and crime can differ across years even when PADO is absent. Appendix Figure A3 shows the same contrast at the more aggregated circuit-calendar scale.

5.6 Heterogeneity by baseline risk and common support

The benchmark follows treated patrol segments. The support exercise instead uses 200-by-200-meter cells because the 2015 robbery-risk measure R15 is defined at that level. It asks whether structural zeros or the highest-risk cells account for the decline. Table 4 re-estimates the controlled core-versus-pre contrast on five pre-specified cell populations: all 827 treated cells; the 753 with positive 2015 risk; the 657 whose risk value also occurs among never-treated cells; and the 674 and 738 cells at or below 14 and 19 robberies in 2015, that is, below or near the selection threshold. In levels, the decline becomes smaller as the highest-risk cells are removed, from -0.555 to -0.351, but so does the matched-2015 robbery mean. As a share of each population's own mean, the point estimate is 29 to 38 percent everywhere, and if anything larger in the cooler cells, while the police first stage changes little. The negative sign is not generated by zero-risk cells, and the drop is not an artifact of a few extreme streets. The pattern is consistent with a roughly proportional effect, but can also combine genuine treatment heterogeneity and differential reversion to the mean, so we do not claim a uniform effect across baseline risk.

Table 4: The robbery decline persists across common-support populations

	Robberies			Police presence	
Cell-level population (cells)	Estimate	2015 mean	% of mean	Estimate	2015 mean
Full support (827)	−0.555*** (0.093)	1.903	−29.2	0.221*** (0.019)	0.187
Positive 2015 risk, R15 > 0 (753)	−0.586*** (0.093)	2.027	−28.9	0.225*** (0.019)	0.191
Exact common support (657)	−0.431*** (0.074)	1.296	−33.2	0.217*** (0.019)	0.177
Trim R15 ≤ 14 (674)	−0.351*** (0.066)	0.932	−37.6	0.211*** (0.020)	0.162
Trim R15 ≤ 19 (738)	−0.407*** (0.085)	1.194	−34.1	0.212*** (0.020)	0.175

Notes. Each row re-estimates the core-versus-pre contrast of Table 2 on a subset of treated 200-by-200-meter cells; R15 is the cell's 2015 robbery count. *Full support*: all treated cells. *Positive baseline risk*: cells with at least one robbery in 2015. *Exact common support*: cells whose R15 value also occurs among never-treated cells. *Trims*: cells with at most 14 or 19 robberies in 2015, at or below the selection threshold. Cell-by-episode and year-month-by-clock-hour fixed effects. Robberies per 1,000 cell-hours; police hours per cell-hour; percentage is the estimate over its 2015 mean. Standard errors clustered by circuit (27 to 29). Rows are different populations, not a test of equal coefficients. *** $p < 0.01$.

6 Deterrence or displacement?

A reduction in the patrolled place-hours could be undone if offenders wait until the patrol leaves, walk to the next street, or switch to another offense. We follow robbery into later hours, into other offense types, and into neighboring cells.

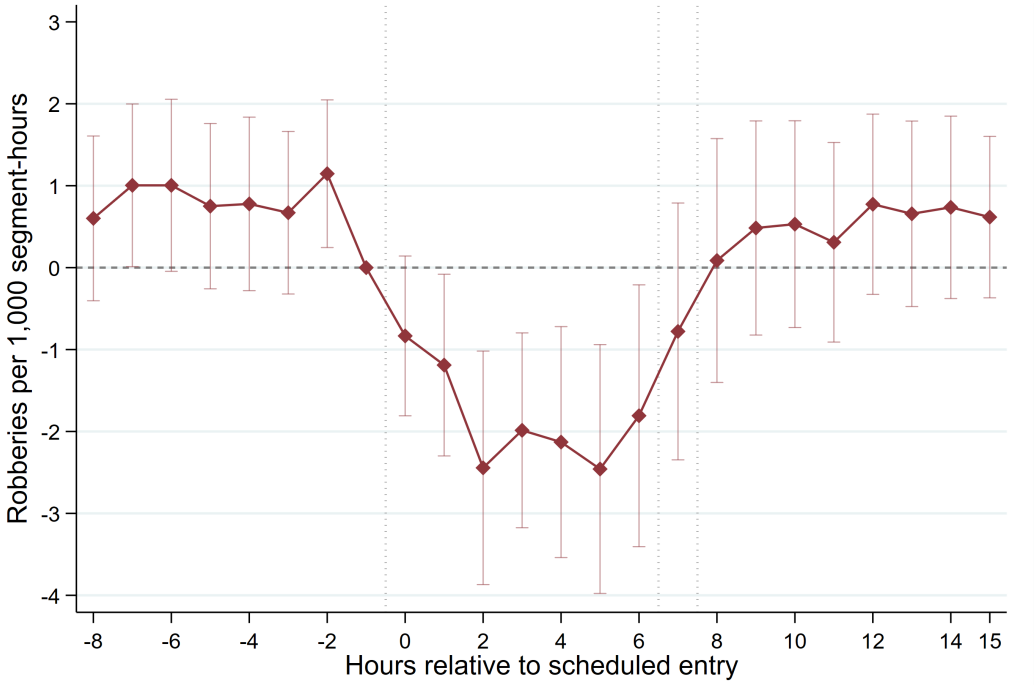
6.1 Temporal displacement

The clean temporal sample requires eight inactive hours before entry and eight inactive hours after the transition. We estimate the same difference specification with phases for the early pre-period, the seven core hours, the transition hour, and two four-hour post blocks, with $k = -4, \ldots, -1$ as the omitted baseline. The primary temporal displacement estimand is the duration-weighted sum of the two post-block coefficients; we also report the seven-core-plus-eight-post total and the compensation share, the post total divided by the absolute core total, with uncertainty from the same segment-clustered model.

Figure 5 and Panel A of Table 6 report the results. In the clean-window population the core robbery coefficient is −2.033 per 1,000 segment-hours against a matched-2015 mean of 5.187, a 39 percent reduction; the transition estimate is −0.912 against a mean of 4.400. The eight post hours sum to +1.403 per 1,000 segment-episodes (standard error 2.688) against a 2015 total of 19.870, and the combined seven-core-plus-eight-post estimate is −12.825 (standard error 4.545) against a 2015 total of 56.180. Over the fifteen-hour horizon robbery is clearly down. The implied compensation share is 0.099: about 10 percent of the core reduction reappears later, and the estimate is not distinguishable from zero. The post coefficient is a net change, however: continued deterrence after officers leave could offset some retiming into the same hours, so a zero estimate does not prove that no offender rescheduled an offense, and the upper end of the compensation interval allows the post

period to offset almost half of the estimated core reduction.

Figure 5: Robbery dynamics before, during, and after clean PADO episodes.



Notes. Hourly matched-2015 robbery coefficients for episodes with eight inactive hours before entry and eight after the transition; $k = -1$ is omitted. Segment-by-episode and calendar year-month-by-clock-hour effects; bars are 95 percent confidence intervals clustered by longitudinal segment.

6.2 Substitution across offense types

A robbery-only result could be misleading if offending changes composition. Within the same treated segments and scheduled hours, we estimate theft and the linear sum of robbery and theft in the same matched-2015 design as Table 2, together with robbery plus theft plus homicide (Table 5). Theft does not rise: the point estimate is −0.492 (standard error 0.337), a 7 percent fall relative to the 2015 mean of 6.906, not significant at conventional levels. The robbery-plus-theft coefficient is −1.846 and its confidence interval remains below zero. Adding homicide barely changes the result because homicide is 0.3 percent of the three recorded categories. We find no substitution into theft; if anything, deterrence spills over to other offenses. The theft interval allows an increase of up to 0.171, which would offset about an eighth of the robbery reduction, so small compositional changes and substitution toward unmeasured offenses cannot be excluded.

Table 5: Direct benchmark and substitution across offense types

Outcome, core hours	Estimate	Standard error	95% confidence interval	2015 mean
Robberies (primary)	−1.354***	0.281	[−1.907, −0.801]	4.479
Theft	−0.492	0.337	[−1.156, 0.171]	6.906
Robbery plus theft	−1.846***	0.417	[−2.669, −1.023]	11.385
Robbery plus theft plus homicide	−1.838***	0.419	[−2.664, −1.012]	11.422

Notes. Events per 1,000 segment-hours during the seven core patrol hours relative to four pre-entry hours, net of the same contrast in matched 2015 observations. All rows use the benchmark sample, fixed effects, and segment-clustered inference. *** $p < 0.01$.

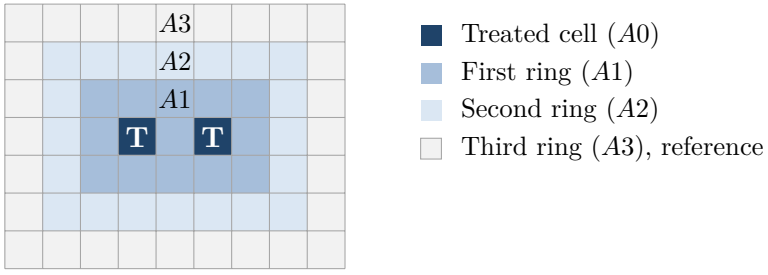
6.3 Local spatial spillovers

The spatial analysis focuses on the six synchronized launch dates, when a whole new program version switched on under a common schedule. These dates permit a transparent comparison between newly treated cells and successive neighboring rings; later redesigns mix isolated starts, overlapping patrol maps, and previously exposed areas, so adding them would increase observations while changing the meaning of the comparison. For the common launch schedule we construct mutually exclusive queen-contiguity rings around the union of treated cells: treated cells ($A0$), first neighbors ($A1$), second neighbors ($A2$), and third neighbors ($A3$), which serve as the spatial reference (Figure 6). The launch map contains 257 treated cells and, after deduplication of overlapping fronts, 857, 924, and 869 cells in the three rings. We estimate

$$\Delta Y_{idk} = \alpha_i + \lambda_{dk} + \sum_{p \neq B} \sum_{r=0}^{2} \beta_{rp} 1\{r(i) = r\} 1\{k \in p\} + u_{idk},$$

where i is a cell, d a launch date, and $k = -4, \ldots, -1$ the temporal baseline; cell fixed effects absorb time-invariant ring differences and date-by-relative-hour effects absorb common launch-day shocks. Spatial inference uses a Conley–Bartlett covariance estimator with a 1 km cutoff, with scores first aggregated within cell.

Figure 6: Spatial groups in the local displacement design.



Notes. Stylized 200-by-200-meter grid. Rings are shortest queen-contiguity distances from the union of treated cells; overlapping fronts are counted once, so ring size need not increase monotonically.

Panel B of Table 6 reports the core-hour contrasts. In treated cells, police presence is 0.209

hours per cell-hour higher than in $A3$, double the matched-2015 treated-cell mean of 0.198, and robbery is 1.715 per 1,000 cell-hours lower, 43 percent of its mean of 3.984. Presence also diffuses modestly into the first ring (0.027 on a mean of 0.098). The robbery contrast for the first ring is −0.057 on a mean of 0.778: no rise next door. The first-minus-second-ring contrast is −0.402. Neither point estimate displays a positive first-ring gradient.

The test is blunt, however. The 80 percent minimum detectable effects are 1.05 for $A1 - A3$ and 1.26 for $A1 - A2$, so robbery in the neighboring ring would have had to more than double before we could detect it. Moreover, the first ring contains many more cells than the treated area, so a modest increase per neighboring cell could offset a much larger decrease per treated cell in total robberies. If the robberies prevented in the 257 treated cells moved next door and spread over the 857 first-ring cells, each ring cell would rise by about 0.5; the first-ring interval permits even this full local compensation of the treated-cell point estimate. For the sum of robbery and theft, the first-ring estimate points the wrong way: +0.901 (standard error 0.637; $p = 0.157$) against a mean of 1.834, a 49 percent rise that is not significant but that we report rather than hide, since theft is the more mobile margin. We therefore detect no positive first-ring gradient in robbery, but cannot rule out an economically important local offset or a more distant response. The evidence is inconsistent with a large, sharply localized displacement of robbery; it is not evidence of its absence.

Table 6: Temporal compensation and local spatial spillovers

Panel A. Temporal compensation, clean 24-hour sample

Estimand	Estimate	Standard error	2015 mean
Core patrol hours, hourly effect	−2.033***	0.390	5.187
Transition hour, hourly effect	−0.912	0.579	4.400
Eight post-patrol hours, total	1.403	2.688	19.870
Seven core plus eight post hours, total	−12.825***	4.545	56.180
Post/core compensation share	0.099	0.197	–

Panel B. Local spatial gradients during the six launch dates, core hours

Contrast	Police presence	Robberies	Robbery plus theft
Treated cells, $A0 - A3$	0.209*** (0.021)	−1.715 (1.044)	−0.896 (1.671)
2015 mean, treated cells	0.198	3.984	6.856
First ring, $A1 - A3$	0.027*** (0.005)	−0.057 (0.375)	0.901 (0.637)
2015 mean, first ring	0.098	0.778	1.834
Second ring, $A2 - A3$	0.004 (0.003)	0.346 (0.333)	1.106 (0.728)
First versus second ring, $A1 - A2$	0.023*** (0.004)	−0.402 (0.449)	−0.204 (0.741)

Notes. Panel A uses the clean 24-hour segment sample and segment-clustered inference. The first two means are hourly; the next two are totals over the hours named in the row. Panel B follows 2,907 cells over the six common launch dates with cell and date-by-relative-hour fixed effects; standard errors are Conley–Bartlett with a 1 km cutoff. Police presence is in hours per cell-hour and crime outcomes in events per 1,000 cell-hours. *** $p < 0.01$.

7 Conclusion

A schedule can move existing officers to the right place *and* the right hour. In PADO, realized police presence more than doubled in the targeted street segments exactly when planned, and robbery in those place-hours fell by about 30 percent. The implied elasticity of robbery with respect to localized patrol exposure is modest, -0.26 to -0.48: the effect is big because the dose is big. Offenders respond to police with a lag. The reduction is about half as large in the first patrol hour as in a mid-shift hour, it then plateaus, and it is still present in the transition hour after officers start returning to base. Read against a precise zero from ten extra minutes of patrol a day in Essex and large losses from withdrawn patrol in Dallas, these results suggest that what deters is presence that is large and sustained enough to be noticed.

The causal interpretation remains conditional. Matching the same geography, weekday, and clock hour in 2015, comparing scheduled with immediately prior hours, and applying the eventual segment-hour schedule to pre-program weeks all make a purely mechanical timing-reversion account less plausible. They do not make timing random. PADO deliberately targeted high-risk place-hours, and later schedules could respond endogenously. The displacement results are bounded in the same way. We detect no robbery rebound during the eight hours after patrol, no rise in theft, and no positive first-ring gradient at launch. Yet the temporal interval allows partial retiming and the spatial tests cannot rule out substantial local compensation or more distant displacement. The evidence narrows the plausible offsets; it does not establish that displacement is zero.

The policy implication is that patrol allocation has a time dimension. Departments that identify high-crime streets but ignore their within-day risk profile may leave deterrence gains unrealized. PADO shows that an operational schedule can materially reallocate realized presence toward high-risk hours, and that robbery responds within the shift. Whether other departments obtain the same return will depend on the initial mismatch, officer compliance, offender adaptation, and the mobility of criminal opportunities. We stop short of a citywide welfare calculation: the direct, temporal, and spatial modules use different populations and weights by design, and a net effect would require a common geographic population, temporal horizon, exposure rule, and valuation of displaced offenses.

References

Ariel, B. and Partridge, H. (2017). Predictable policing: Measuring the crime control benefits of hotspots policing at bus stops. *Journal of Quantitative Criminology*, 33:809–833.

Ariel, B., Sherman, L. W., and Newton, M. (2020). Testing hot-spots police patrols against no-treatment controls: Temporal and spatial deterrence effects in the london underground experiment. *Criminology*, 58(1):101–128.

Becker, G. S. (1968). Crime and punishment: An economic approach. *Journal of Political Economy*, 76(2):169–217.

Blanes i Vidal, J. and Mastrobuoni, G. (2018). Police patrols and crime. Technical Report 11393, IZA Institute of Labor Economics. Revised May 2025.

Blattman, C., Green, D. P., Ortega, D., and Tobón, S. (2021). Place-based interventions at scale: The direct and spillover effects of policing and city services on crime. *Journal of the European Economic Association*, 19(4):2022–2051.

Braga, A. A., Turchan, B., Papachristos, A. V., and Hureau, D. M. (2019). Hot spots policing of small geographic areas effects on crime. *Campbell Systematic Reviews*, 15:e1046.

Cabrera, J. M., Cid, A., and Veneri, F. (2022). Hot spots, patrolling intensity, and robberies: Lessons from a three-year program in uruguay. Technical Report 2202, Universidad de Montevideo.

Chainey, S. P., Serrano-Berthet, R., and Veneri, F. (2021). The impact of a hot spot policing program in montevideo, uruguay: An evaluation using a quasi-experimental difference-in-difference negative binomial approach. *Police Practice and Research*, 22(5):1541–1556.

Chalfin, A. (2025). Investments in policing and community safety. *Annual Review of Criminology*, 8:403–429.

Chalfin, A. and McCrary, J. (2018). Are u.s. cities underpoliced? theory and evidence. *Review of Economics and Statistics*, 100(1):167–186.

Di Tella, R. and Schargrodsky, E. (2004). Do police reduce crime? estimates using the allocation of police forces after a terrorist attack. *American Economic Review*, 94(1):115–133.

Draca, M., Machin, S., and Witt, R. (2011). Panic on the streets of london: Police, crime, and the july 2005 terror attacks. *American Economic Review*, 101(5):2157–2181.

Giulietti, C., McConnell, B., and Zenou, Y. (2025). Beyond hot spots: Enhancing police effectiveness by incorporating a spatial network approach. Technical Report 20065, Centre for Economic Policy Research.

Inter-American Development Bank and Ministry of Interior of Uruguay (2018). *¿Cómo evitar el delito urbano? El Programa de Alta Dedicación Operativa en la nueva Policía uruguaya*. Inter-American Development Bank and Ministry of Interior of Uruguay, Washington, DC and Montevideo.

Koper, C. S. (1995). Just enough police presence: Reducing crime and disorderly behavior by optimizing patrol time in crime hot spots. *Justice Quarterly*, 12(4):649–672.

Machin, S. and Marie, O. (2011). Crime and police resources: The street crime initiative. *Journal of the European Economic Association*, 9(4):678–701.

Mello, S. (2019). More cops, less crime. *Journal of Public Economics*, 172:174–200.

Nagin, D. S., Solow, R. M., and Lum, C. (2015). Deterrence, criminal opportunities, and police. *Criminology*, 53(1):74–100.

Sherman, L. W. and Weisburd, D. (1995). General deterrent effects of police patrol in crime "hot spots": A randomized, controlled trial. *Justice Quarterly*, 12(4):625–648.

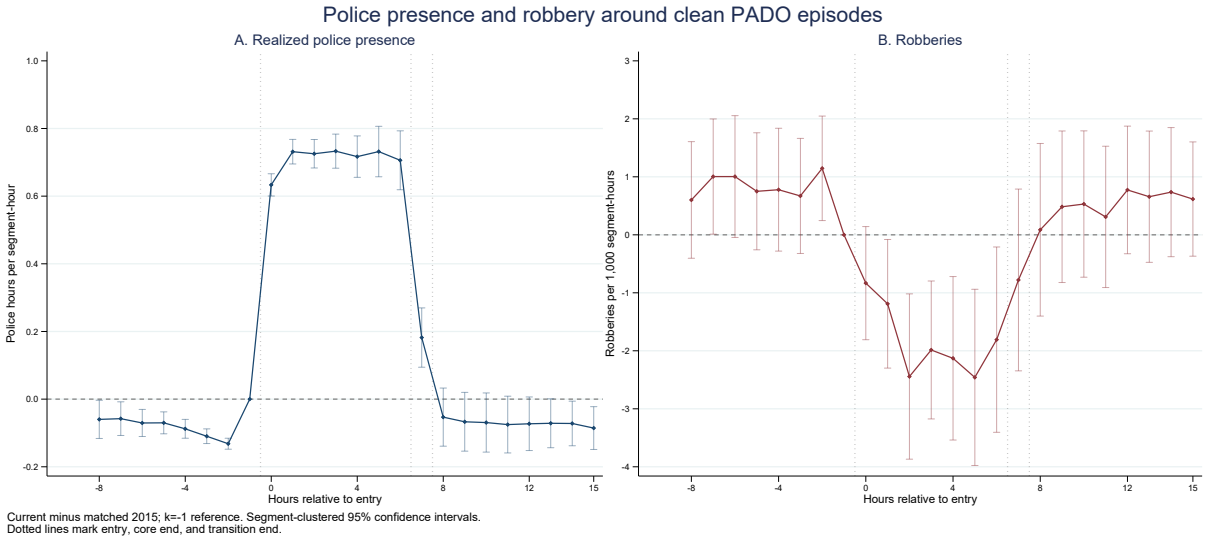
Vandeviver, C., Dau, P. M., Bland, M., Van De Walle, S., Dewinter, M., Witlox, F., and Vander Beken, T. (2026). Short infrequent police stops maximize crime-free time at street segments: Revisiting the koper curve to establish optimal police dosage. *European Journal of Criminology*, 23(4):655–678.

Weisburd, S. (2021). Police presence, rapid response rates, and crime prevention. *Review of Economics and Statistics*, 103(2):280–293.

Williams, S. and Coupe, T. (2017). Frequency vs. length of hot spots patrols: A randomised controlled trial. *Cambridge Journal of Evidence-Based Policing*, 1:5–21.

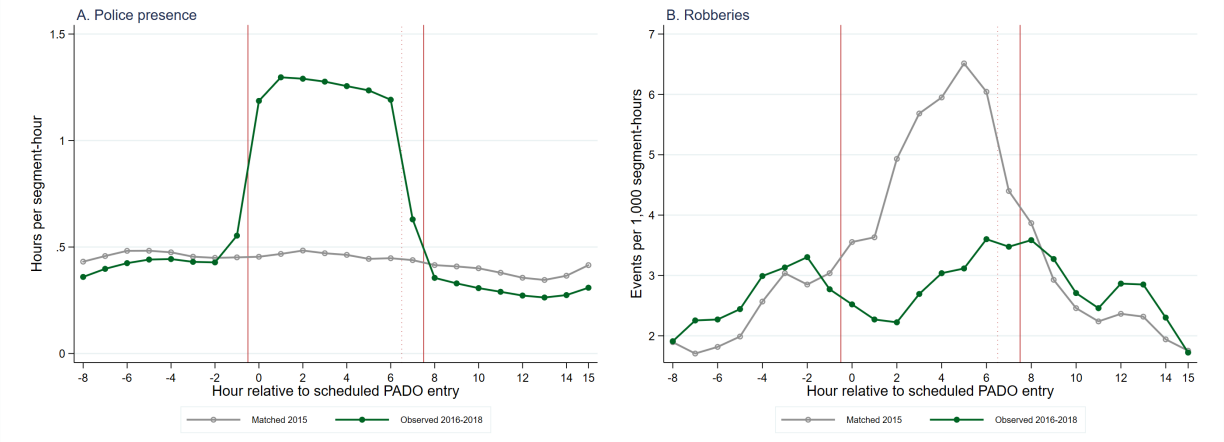
Appendix

Figure A1: Police presence and robbery around clean 24-hour PADO episodes.



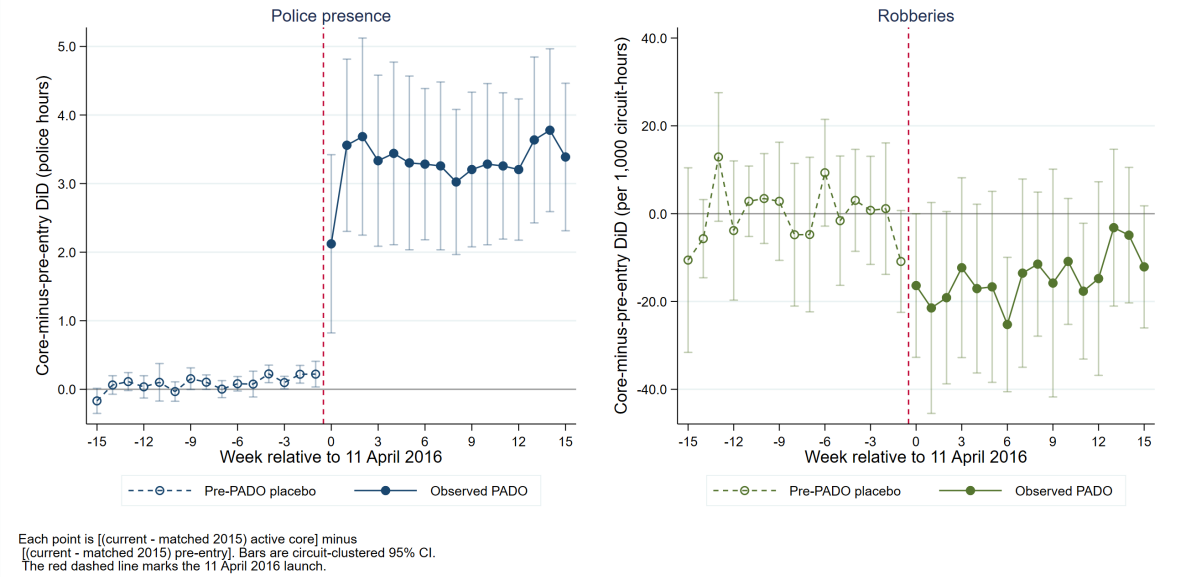
Notes. Current minus matched 2015; $k = -1$ reference; 95 percent segment-clustered confidence intervals. Dotted lines mark entry, the end of the core, and the end of the transition hour. Episodes with eight inactive hours before entry and after transition: 13,149 episodes, 63,866 segment-episode pairs. Appendix Figure A2 shows the same sample in levels.

Figure A2: Outcome levels around clean 24-hour PADO episodes.



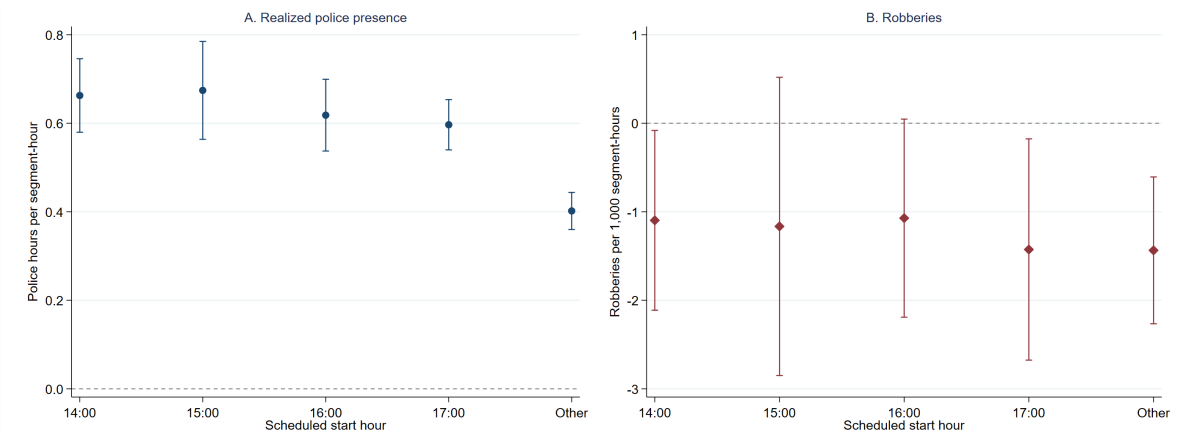
Notes. Observed 2016–2018 outcomes and the matched 2015 profile for the clean temporal-displacement population, without an adjusted counterfactual or uncertainty bands.

Figure A3: Full-schedule timing placebo around the PADO launch, circuit-calendar scale.



Notes. Each point is the core-minus-pre difference in the current year relative to the matched 2015 calendar. Open points project the eventual schedule distribution into pre-PADO dates; solid points use actual schedules after 11 April 2016. Bars are circuit-clustered 95 percent confidence intervals. The main-text placebo asks the same question at the segment-hour resolution of the benchmark.

Figure A4: Core-hours estimates by scheduled start time.



Notes. Estimates and 95 percent confidence intervals from one joint model with month-specific clock-hour controls. "Other" includes 18:00 and all hours outside 14:00–17:00. Matched-2015 core robbery means are 4.193, 5.098, 5.045, 5.678, and 3.443 for the 14:00, 15:00, 16:00, 17:00, and residual groups.

Table A1: Police–crime elasticities in the literature and in PADO

Study and source of variation	Elasticity
Chalfin and McCrary (2018): robbery, city force size	−0.56
Mello (2019): violent crime, force size	−1.3
Di Tella and Schargrodsky (2004); Draca et al. (2011): static deployment after terror attacks	−0.3 to −0.4
Weisburd (2021): Dallas, car patrol withdrawn from beat	−0.9 violent / −0.6 property
Blanes i Vidal and Mastrobuoni (2018): Essex, ten extra minutes a day of foot patrol	≈ 0
PADO, hourly (this paper)	−0.26 / −0.48

Notes. Force-size elasticities (officers employed) and patrol-exposure elasticities (officers present) are not directly comparable. Di Tella and Schargrodsky (2004) do not observe police presence; their elasticity assumes the increase in officers per block. The earlier weekly evaluation of PADO (Cabrera et al., 2022) reports −0.47.

Table A2: Benchmark results under alternative inference rules

	Police presence		Robberies	
Rule	Estimate	Standard error	Estimate	Standard error
Segment CR1 (206 clusters), main	0.517***	0.018	−1.354***	0.281
Segment × episode CR1	0.517***	0.019	−1.354***	0.278
Circuit CR1 (29 clusters)	0.517***	0.048	−1.354***	0.214
Circuit CV3	0.517***	0.068	−1.354***	0.275

Notes. All rows use 1,503,280 observations and the benchmark core-hours contrast. CR1 denotes the conventional cluster-robust calculation; CV3 is a leave-one-circuit-out correction for the 29 circuit codes. *** $p < 0.01$.

Table A3: Implied police–robbery elasticity under alternative normalizations

Normalization	Estimate	Standard error	95% confidence interval
Matched-2015 baseline scale	−0.257	0.052	[−0.360, −0.154]
Midpoint arc	−0.480	0.116	[−0.708, −0.252]

Notes. Both rows use the 15,352 benchmark episodes, seven core hours, fixed effects, and segment-level clustering of Table 2. Uncertainty is calculated from a stacked system so that the covariance between the robbery and police-presence estimates enters the ratio. These are implied elasticities of localized patrol exposure and inherit the benchmark timing assumption.