Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Signatures of the urban calendar in emergency medical services

Chris Lim

chrislim@asu.edu

Arizona State University

Ran Wei

School of Public Policy, University of California Riverside

Jonathan Race

Department of Pediatrics, School of Medicine, University of Utah

Honghyok Kim

School of Public Health, University of Illinois Chicago

Xiaojiang Li

University of Pennsylvania

Lynn Gerald

Office of Population Health Sciences, University of Illinois Chicago

Chen Chen

Department of Population Health Sciences, University of Utah

N. Clay Mann

Department of Pediatrics, School of Medicine, University of Utah

Article

Keywords:

Posted Date: November 19th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11760581/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: There is **NO** Competing Interest.

Signatures of the urban calendar in emergency medical services

Chris C. Lim[1,2]*, Ran Wei[3], Jonathan A. Race[4], Honghyok Kim[5], Xiaojiang Li[6], Lynn Gerald[7], Chen Chen[8], N. Clay Mann[4]

[1]School of Technology for Public Health, Arizona State University, Phoenix, AZ, USA; [2]Mel & Enid Zuckerman College of Public Health, The University of Arizona, Tucson, AZ, USA; [3]School of Public Policy, University of California Riverside, Riverside, CA, USA; [4]Department of Pediatrics, School of Medicine, University of Utah, Salt Lake City, UT, USA; [5]School of Public Health, University of Illinois Chicago, Chicago, IL, USA; [6]Department of City and Regional Planning, University of Pennsylvania, Philadelphia, PA, USA; [7]Office of Population Health Sciences, University of Illinois Chicago, Chicago, IL, USA; [8]Department of Population Health Sciences, University of Utah, Salt Lake City, UT, USA

*Corresponding author: Chris C. Lim (chrislim@asu.edu)

Abstract

Cities are organized by schedules as well as streets. Street permits, benefit calendars and holidays create recurring pulses that may change emergency demand and how long help takes to arrive. In 16.5 million New York City emergency medical services records (2012-2023), we measured call volumes, types and response times within each pulse. During active marathon operations, call-to-scene time rose 152 seconds more near the route than far from it. Dispatcher-coded drug or poisoning calls rose 8.6% on each month's days 1-3, and July 4 fireworks windows had 233% more burn calls. Each schedule had its own geography: marathon delays followed the route, not neighbourhood characteristics, whereas the early-month increase was modified by them. Estimated one at a time, some pulses appeared larger than they were. Set in advance, some by the city, these schedules are a recurring pressure on emergency care that cities can plan for as one calendar, street by street and neighbourhood by neighbourhood.

Cities are increasingly sensed through the data their residents and infrastructures generate. Mobile-phone, transit and service-request records have traced daily and weekly urban rhythms[1–3], and the mix of requests a neighbourhood files with the city's 311 service can characterize a place[4]. These layers seldom record acute health needs, or whether public services can reach the people who have them. A few studies have brought ambulance calls into this work, as neighbourhood rhythms and as predicted volumes[5,6], but the time it takes help to arrive has rarely been measured the same way.

Emergency medical services (EMS) dispatch records add a health-system layer that records both the need and the response. Each incident is time-stamped when it is created in the dispatch system, usually soon after it is recognized. The records can register events lasting only hours, including incidents that never lead to hospital admission or death, and they record how long patients wait for responders[7–10]. Like other urban sensors, they are filtered: incidence, recognition, help-seeking, service availability and dispatcher coding all shape what enters dispatch, so they measure demand visible to EMS rather than underlying morbidity.

Cities are organized by schedules as well as streets. Benefit calendars set payment dates, holidays synchronize gatherings and changes in service availability, and permits reserve road space for public events. Following scholarship on social time and infrastructural rhythms[11–13], we treat schedules as temporal urban infrastructure and the recurring periods they mark as urban pulses. Payment-timing studies associate disbursement calendars with short-run changes in consumption, drug use, crime and mortality[14–19], and earlier emergency studies found more ambulance runs and 911 calls after welfare cheques were issued[20,21]. Research on marathons and mass gatherings documents clinical and operational disruption and higher EMS call volume[22–24]. Across these literatures, health evidence comes mostly from death, hospital and emergency-department records, and EMS studies have more often counted calls than timed responses. The few that timed responses each followed a single kind of event, and their results diverged[23,25,26]. Because each literature mostly examines one kind of schedule, different schedules are seldom compared within the same city, and it is unclear whether a schedule's effect falls on how many people call, on how long they wait, or on both.

EMS demand also has its own calendar. Ambulance calls vary by weekday, season and year, and their daily rhythms differ between neighbourhoods[5]. A holiday can therefore look unusual only because it behaves like a Sunday, and a fixed date can fall within another pulse, as July 4 falls in the days after each month's start. Taken one at a time and through call counts alone, a schedule's size can be misjudged and its cost in waiting missed.

We examine three kinds of schedule in 16.5 million publicly available New York City EMS records from 2012 through 2023[27]: permits and security plans that close streets, the first days of each month, when benefit calendars concentrate many payments, and holidays. Each is fixed in advance in public records, so their pulses can be separated and compared within one system. For each, we measure demand visible to EMS, the incidents entering dispatch and their types, and pre-arrival access: call-to-scene time, from incident creation to scene arrival. The street schedules produce street-operation pulses; the benefit and holiday schedules produce calendar pulses. We call the combination of call patterns and response times a schedule leaves its signature. Later New York City data and openly published dispatch records from three other US cities test which patterns recur.

Because these pulses share one calendar, a pulse measured on its own can include part of another. A joint calendar model therefore places every pulse on one timeline, beneath regular weekly and seasonal variation, and

neighbourhood models ask which places carry each pulse. Where a pulse falls may depend on what sets it: the route of an operation, or the residents whose days a calendar reorganizes. To compare schedules that change demand with those that change access, we also sum pre-arrival waiting across incidents, multiplying calls by their mean call-to-scene time; the change in this total separates more calls from slower responses. Measured together rather than one at a time, and against the right reference days, some holiday excesses shrank or reversed, so a city's schedules are better planned for as one calendar.

Results

Access to patients slowed during large street operations

Call-to-scene delays concentrated in active marathon operations. Across the 10 years in which the race was held, annual route-adjacent estimates, weighting each year equally, averaged +152 seconds (95% CI 93-210; matched-date permutation $P < 0.001$; Fig. 1a,c) against a control-Sunday mean of 476 seconds, and were positive every year (Supplementary Fig. S5a; Supplementary Table S23). Marathon Sundays also had 15.2% more route-adjacent incidents (95% CI 5.4-26.0%, excluding event-standby records; Supplementary Table S12). The delay was not confined to calls the event could have generated: call types unlikely to come from runners or spectators, such as psychiatric, obstetric and stroke calls, were delayed by 118 seconds (95% CI 58-179), and life-threatening calls by 151 seconds (108-194; Supplementary Table S48).

Both stages of the ambulance response slowed, travel more consistently than dispatch. In the incidents behind the call-to-scene estimate, dispatch took 68 seconds longer (95% CI 8-129) and travel 83 seconds longer (53-113); in records with a complete call-to-hospital interval, the dispatch estimate only just excluded zero (+37 seconds, 95% CI 0.1-74.0; cluster-robust $P = 0.055$; Extended Data Table 2). Fire companies on non-medical incidents slowed less, by 25 seconds (95% CI 4-46; Supplementary Table S27).

Delays also occurred during United Nations General Assembly weeks. Annual call-to-scene estimates in the three UN-adjacent ZIPs averaged +113 seconds across 11 in-person years (95% CI 71-154) and were positive in every year, whereas the 2020 virtual session showed no detected association (-28 seconds, -88 to 32; Supplementary Tables S11 and S23).

To test whether the delay extended to other street events, we classified New York's 2024-2025 permitted street events by their street closures before examining EMS outcomes. Parades and races that closed full streets across two or more police precincts were associated with 37.4 seconds longer precinct-hour call-to-scene time (95% CI 15.8-59.1), whereas single-precinct full closures and less-restrictive permits showed no detected association and placebo windows before and after the events showed none ($P = 0.39$; Fig. 1c; Supplementary Tables S16 and S23). The association was concentrated in Manhattan. Excluding the 13 federal-holiday parade dates left +30.2 seconds (10.5-49.9), but outside Manhattan it was +12.7 seconds (-12.0 to 37.3), and without both it failed a prespecified generalization rule (+9.2 seconds, -14.8 to 33.3; Supplementary Table S49). Passenger-trip records independently showed longer trips during permitted events. Across 263 occurrences on 166 dates, trip-duration contrasts averaged +54 seconds (95% interval 43-65) and correlated with EMS call-to-scene contrasts at $r = 0.289$ (0.155-0.403; Table 1; Supplementary Fig. S4; Supplementary Tables S26 and S28).

Drug or poisoning calls concentrated in the early-month pulse

Dispatcher-coded drug or poisoning calls clustered in the first three days of each month. Across 396 matched case-day strata (January excluded), days 1-3 had 8.6% more such calls than same-weekday control days later in the month (95% CI 7.4-9.8%, P < 0.001). Calls rose on day 1, peaked on day 2 and decayed over about a week (Fig. 2b).

The drug or poisoning share of all calls also increased: the adjusted odds of this category versus all other calls were 8.1% higher (OR 1.081, 95% CI 1.068-1.094), an odds ratio rather than a percentage-point change in the proportion of calls (Extended Data Table 3). All EMS calls combined rose only 1.2% (0.6-1.8%), so the increase was not a uniform rise in call volume.

The concentration appeared in every borough, income quartile and four-year period, and was largest in summer and smallest but positive in winter (December and February; Extended Data Table 4; Supplementary Fig. S2 and Tables S1, S3, S7 and S18). Advancing Supplemental Security Income (SSI) payments did not attenuate the days 1-3 contrast, although calls rose from the advanced payment date to the month boundary (incidence rate ratio [IRR] 1.070, 1.037-1.104).

Assault, burn and drug or poisoning calls rose on holidays

Category-specific call counts increased sharply on July 4 (Fig. 3). In the July 4 fireworks window (18:00-05:59, 12 annual event windows), burn calls were 233% higher than in matched hours 1 and 2 weeks before and after (95% CI 144-356%), assault calls were 133% higher (92-183%) and drug or poisoning calls were 34% higher (21-49%), partly reflecting the date's place in the month (see below). Cardiac-arrest counts showed no detected increase across the window (-3 to 16%) but rose 21% in an exploratory post-midnight window (4-42%; Supplementary Table S2).

On Memorial and Labor Day, against matched weekdays across 24 holiday occurrences, assault calls rose 39% (25-55%) and drug or poisoning calls 16% (10-23%), cardiac-arrest counts showed no detected change (-7 to 7%; Extended Data Table 5), and mental-health calls fell 16% (13-18%), most during business hours (Extended Data Table 6). Across the full holiday day in 2012-2023, responses were faster: call-to-scene times were shorter on July 4 (-68 seconds, -94 to -41) and on Memorial and Labor Day (-89 seconds, -114 to -65; Supplementary Tables S12 and S25). Within the July 4 fireworks window itself, however, call-to-scene time was 88 seconds longer (49-124; Fig. 6; Supplementary Table S46).

Each schedule had its own geography

Across 170 ZIPs, each pulse clustered in space (Moran's I 0.38 for the early month and 0.29 for Memorial and Labor Day among the 168 ZIPs with neighbours; both P = 0.001; Fig. 2a), but neighbourhoods with large excesses in one pulse did not clearly have large excesses in others (Supplementary Table S32). The early-month and Memorial and Labor Day maps were nonetheless reproducible: each reproduced itself between odd and even years (estimated correlations between underlying ZIP-level effects 0.71 and 0.70) more closely than it resembled the other pulses' maps (0.01 and 0.21). The July 4 map, with six occurrences in each half, was too imprecise to test (Fig. 4b-d; Supplementary Table S50).

The early-month increase fell with median income (-1.9% per standard deviation), and its shape over days 1-9 differed across income quartiles (Fig. 2c). With the share of non-Hispanic Black residents in the same model, the

income association was no longer evident (0.0%, -1.2 to 1.2), and the increase was 3.4% larger per standard deviation of that share (about 24 percentage points; 2.3-4.5%; Supplementary Fig. S5e). The calendar pulses, including New Year's Day, did not share one neighbourhood profile ($P < 0.001$; Fig. 4e; Supplementary Table S50). Income, racial composition and public housing together accounted for 72% of the early-month map's variation across ZIPs but none of the July 4 map's, which was smaller in denser ZIPs.

Marathon delays concentrated near the route. The +152-second contrast combined a route-adjacent increase of 195 seconds (120-269) with a smaller increase in distant comparison ZIPs (+43 seconds, 3-83), and changes declined with distance from the route (Fig. 1b; Supplementary Table S34). Every held marathon fell on the Sunday when daylight saving time ends, and the distant change cannot be separated from that clock change or from units redeployed for the race. The route contrast did not depend on the clock change: against standard-time Sundays alone it was +157 seconds (97-216), and the same design gave +1 second (-21 to 24) on the 12 Sundays when daylight saving time began (Supplementary Table S49). The delay also followed the route rather than the neighbourhood, whereas the calendar maps showed no detected association with distance from the course ($P = 0.39$). In an incident-weighted model holding income, racial composition, public housing, density and hospital access fixed, route-adjacent ZIPs still slowed by 122 seconds more than distant ones (95% CI 33-210), and the intervals for income, racial composition and public housing stayed within 40 seconds per standard deviation (Fig. 4a,e; Supplementary Table S50).

Overlapping schedules altered estimated holiday effects

A joint model placed all calendar pulses on one timeline of EMS demand, with weekday and lagged-temperature terms and January included (Fig. 5; Supplementary Table S31), and the early-month increase extended beyond days 1-3. In a separate overlap test using the matched designs, day 4 still ran 6.5% above mid-month in other months (4.7-8.3%). Because July 4 always falls in this tail, accounting for the early-month component lowered the full-day July 4 drug or poisoning excess from 22.3% to 14.9% (95% CI 6.3-24.1%; Fig. 5b; Supplementary Table S31, Panel F).

Compared with the same month's Sundays instead, holiday drug or poisoning calls were lower: 6.5% lower on Memorial Day (3.5-9.3% lower) and 4.5% lower on Labor Day (0.2-8.5% lower). Much of their rise against weekdays therefore reflected the weekly day-off rhythm. Assault calls diverged. Labor Day, when J'Ouvert and the West Indian American Day Parade fill Eastern Parkway in Central Brooklyn, ran 18.1% above the Sunday pattern (3.4-34.8), as did weekday July 4s (37.6%, 20.5-57.2%), but not Memorial Day (Fig. 5c; Supplementary Table S31).

Street operations and July 4 nights combined more calls with slower responses

Multiplying call counts by mean call-to-scene time gave each schedule's total pre-arrival waiting, an aggregate of waiting hours summed across incidents rather than observed per-patient time (Fig. 6; Supplementary Table S46). In route-adjacent ZIPs, total waiting was 64% higher during active marathon operations (95% interval 41-91), or 28 incident-hours per race (17-43) above a comparison level of about 44; 39% of the excess came from more calls, 46% from slower responses and the remainder from both. In the July 4 fireworks window, total waiting was 41% higher (28-54), as calls rose 20% and call-to-scene time rose 88 seconds (49-124), even though responses across the full holiday day were faster.

Total waiting changed less, or fell, on other schedules. It was 19% higher on United Nations week days (13-25), mostly from slower responses, and showed no detected change on early-month days (+1.8%, -0.3 to 4.4). Against the Sunday pattern, it was 6% lower on Memorial and Labor Day (1-10), from faster responses. Placebo Sundays and the virtual 2020 session showed no excess (Fig. 6; Supplementary Table S46).

Call patterns recurred in later New York data and in other cities

When the prespecified calendar contrasts were reapplied to 2024-2025 data, all seven predicted directions matched, and the two periods' log-IRRs across all 15 schedule-by-endpoint cells had a Spearman correlation of 0.914 (Fig. 7a; Supplementary Fig. S5d). The response-time pattern for calendar events did not fully recur: the full-day July 4 estimate changed sign on its two later occurrences, so a prespecified rule that required both patterns to recur, with an early-month estimate no larger than +30 seconds and negative estimates for both holiday classes, classified the combined signature as not supported (Fig. 7b; Supplementary Table S18). Street-operation delays did recur. Applied unchanged to 2024 and 2025, the marathon design on the corrected course gave +91 and +164 seconds, and the UN design +151 and +66 seconds (Supplementary Table S49). Call-count directions also recurred in two other metropolitan areas, Maricopa County (Phoenix and Glendale) and Cincinnati, where 16 of 18 prespecified positive-direction cells had IRRs above 1 and all eight cardiac-arrest comparison intervals included 1 (Supplementary Table S24). The early-month pattern was directionally consistent there but not detected: all three estimates exceeded 1, and none excluded no change.

Discussion

Urban schedules place recurring pressures on emergency medical services that total call counts alone would miss. Urban sensing has traced the city's rhythms mainly through movement and activity[1,2,28]. Dispatch records add how long the people who called for help waited for it. Near marathon routes and the United Nations, waits were longer in every year the events were held, from 2012 through 2025. Calendar pulses mainly changed which calls reached dispatch, and marathon Sundays also brought more calls near the route. Unlike heat waves or storms, these pressures follow calendars and permits fixed in advance, some issued by the same city that runs the ambulance service.

Each kind of schedule also had its own geography. The marathon delay followed the route and changed little with a neighbourhood's income, racial composition or public housing. Calendar pulses followed neighbourhood character instead, each in its own way. The early-month increase in drug or poisoning calls was larger where more residents were Black, conditional on median income, benefit receipt and borough, and where public-housing shares were larger. With racial composition in the model, the income association was no longer evident. At the ZIP-code level the two are correlated, and Black share may index unmeasured deprivation, drug supply or help-seeking that these data cannot separate; the associations attach to incident locations and do not identify individual risk.

The schedules also overlap, because they share one calendar with each other and with the weekly and seasonal rhythms beneath them. About a third of the apparent full-day July 4 drug or poisoning excess was accounted for by the month cycle, and in the joint model Memorial and Labor Day resembled Sundays; the first is an overlap adjustment, the second also a change of reference day. Adding the pulses that coincide on a date gives the total excess to plan for, and separating them shows how much belongs to the holiday, the month cycle or the day off.

Relation to earlier evidence

Whether street events delay the ambulance's arrival had been inferred rather than measured: earlier work documented longer ambulance transport times on marathon days and higher mortality among older patients hospitalized with acute cardiac conditions, and took similar scene-departure times in that restricted sample to mean ambulances were not delayed in reaching patients[23]. Measured directly across all dispatched calls, call-to-scene time was 152 seconds longer near the route, and the comparison extends to United Nations General Assembly weeks and to permitted street events classified before EMS outcomes were examined. Travel slowed and dispatch probably did too, although its estimate was less stable across samples; if so, the delay did not arise from road closures alone. Fire companies on non-medical incidents, which face the same closures but are not occupied by patient care and transport, slowed by about a sixth as much as ambulances. Patient outcomes were unavailable, leaving the clinical consequences of these waits unmeasured[29,30].

Earlier evidence from emergency services rested on counts of ambulance runs and 911 calls after welfare cheques[20,21]. Twelve years of categorized dispatch records narrow that signal: the excess sits in drug or poisoning calls on days 1-3, returns to the same neighbourhoods in alternate years, and matches the timing of payment-cycle findings for mortality, hospital admissions, overdose and crime[15,17–19]. Its timing is better established than its cause: advancing SSI payments did not attenuate it, and rent and some payroll dates also fall at the month boundary.

Emergency-department surveillance establishes the seasonal peak of fireworks injuries[31,32]; dispatcher codes registered the same burn excess hour by hour on the night itself. The nulls mark the limits of that signal: cardiac-arrest counts showed no detected rise across the July 4 fireworks window or on Memorial and Labor Day, apart from a 21% rise after midnight on July 4, whereas cardiac mortality rises around winter holidays[33].

Recurrence

Call-count directions held in later New York and external-city data, and street-operation delays were positive in every held year. The response-time pattern that failed the prespecified test concerned citywide calendar events: full-day July 4 reversed on two later occurrences, while Memorial and Labor Day remained faster. Because the 2024-2025 test was only partly prospective and its sample small, the reversal shows instability rather than absence of the earlier association. Call magnitudes across the prespecified cells were also about 40% smaller in the later period: agreement in direction does not calibrate a forecast, and the later analyses did not test whether call shares recurred. A scan of other dates found further holiday pulses, notably Halloween assaults and drug or poisoning calls on the eve of Thanksgiving, but no citywide pulse for the religious and lunar observances tested (Supplementary Table S47; Supplementary Note 1).

Limitations

Dispatcher categories record what callers report and call-takers code[27], not clinical diagnoses. The observed mix can change with incidence, recognition, help-seeking, use of other services and coding, and with who is present in a neighbourhood. The drug-or-poisoning category is not opioid-specific, and fewer mental-health-coded calls need not mean less distress.

The event designs cannot separate closure, crowds, staging, policing, dispatch prioritization and incident demand. Exposure was assigned by ZIP centroid, which misclassifies some incidents in large or irregular ZIPs, and fewer response intervals were valid near the route, although assigning one second to every invalid active-window interval still left +104 seconds (61-147; Supplementary Tables S19 and S23). The passenger-trip comparisons cover selected completed trips with unobserved routes and cannot establish that passenger-transport disruption mediated the EMS delay. Estimates from the citywide calendar designs and the ecological neighbourhood and street-event designs are not directly comparable.

Public records did not identify event-deployed ambulances, EMS capacity or individual routes, and density, road-network form and dispatch practice may change access elsewhere. Within New York, the permit association was concentrated in Manhattan, where street grid, event scale, crowds and policing differ from other boroughs. The 2025 data also follow congestion pricing below 60th Street, which shortened ambulance travel by about a minute[10], and a March 2025 FDNY directive to transport patients to the nearest hospital[34]; the permit association was smaller and less certain in 2025 than in 2024 (Supplementary Table S16). The schedules studied here are also examples rather than a complete set; neighbourhood-scale observances, transit service changes, other benefit dates and large outdoor concerts were not tested.

Implications for city policy

New York's government both issues street-event permits and runs the ambulance service. Because local permitted events showed no detected delay and multi-precinct delays were concentrated in Manhattan, an EMS access review could start with the multi-precinct full-street class there, tracking the number and location of long waits event by event. If dispatch slowed as well as travel, as the principal estimates suggest, unit availability matters as well as crossing points.

Calendar pulses call for a different response, because they concentrate in particular neighbourhoods rather than along a route. The early-month pulse was small citywide, about 22 calls a day without a detected change in total waiting, but it recurred, and its neighbourhood map held between alternate years. That pattern favours targeted surveillance and outreach: naloxone distribution on days 1-3 where the pulse has been largest, chosen by observed excess rather than neighbourhood demographics[35], and burn and trauma readiness on July 4 nights[32]. These data do not isolate a payment mechanism, and none of these responses has been evaluated against these pulses.

New York sets some of these schedules and publishes the dispatch records that show their signatures, and it is not alone: the calendar contrasts recurred in public dispatch records from Phoenix, Glendale and Cincinnati. Any city that publishes such records can read its own calendar in them, plan for the demands and waits that accompany it, and measure, neighbourhood by neighbourhood, whether those plans worked.

Methods

Data source and outcome construction

We analysed New York City's public EMS Incident Dispatch Data for 1 January 2012 through 31 December 2023. The frozen raw extract contained 17,490,489 study-period rows before analytic-geography and weather restrictions. Primary calendar analyses used 16,514,279 incidents in the complete 170-ZIP NLDAS-linked panel; incident-level

supporting models used 16,892,512 rows satisfying their required geography and clock fields. The calendar-boundary, holiday, marathon and United Nations General Assembly analyses use this period. As in the source extraction, the calendar panel counts only incidents with a valid response interval, 96.9% of incidents in these ZIPs. Counting all incidents changed no primary estimate's direction or whether its confidence interval excluded 1 (Supplementary Table S22). The 2024-2025 validation counts included incidents without a valid interval.

We used 2024-2025 data for three tests: whether the calendar patterns recurred, under contrasts fixed in a dated protocol before the later counts and response times were extracted; whether the marathon and UN delays persisted, under the designs reapplied unchanged, with one record-count rule amended before any estimate was computed (Supplementary Table S49); and whether delays extended to other permitted street events, with exposures classified before any EMS outcome was examined. The dispatch file contains one row per incident, with intervals for the first responding unit. CAD_INCIDENT_ID identifies records in the frozen raw extract; ZIP-by-timestamp collisions, checked only as a diagnostic, were below 0.03% in every corrected marathon design cell (Supplementary Table S5).

NYC Open Data exposes INCIDENT_DATETIME as a floating civil timestamp without an encoded timezone. We parsed the source text as America/New_York wall time, then converted the localized value to a UTC instant before deriving the local date, hour, weekday and event window. The primary holiday windows and active marathon window do not overlap the repeated fall-back hour; ambiguous records, weather linkage and the daylight-saving audit are described in Supplementary Methods and Supplementary Table S20.

For count-based analyses we aggregated incident records to city-hour, city-day or ZIP-day panels. Endpoint categories (drug or poisoning, assault, mental-health, cardiac-arrest, burns, trauma, unconscious, seizures, and total calls) are derived from FDNY EMS dispatch call types using FINAL_CALL_TYPE, the call type when the incident closes. The source dictionary describes call types as determinations from the caller's information that do not change with the on-scene crew's findings; they are not paramedic impressions or hospital diagnoses (Supplementary Table S21). The categories are mutually exclusive because each record is assigned from one final code. For the marathon analysis we also examined dispatch, travel, scene, transport and total response times. Call-to-scene time runs from incident creation in the dispatch system to the first unit's arrival and excludes 911 call-taking before the incident is created. We interpret EMS counts as joint measures of incident occurrence, recognition, help-seeking, service availability, and dispatcher coding rather than as direct measures of latent morbidity alone.

Endpoint definitions, record accounting and the dated analysis ledger are in Supplementary Tables S21 and S22. The designs compare patterns within one emergency system; effect magnitudes are not directly comparable.

Early-month calendar-boundary case-crossover

For the early-month analysis, drug or poisoning and comparison endpoints were aggregated to the city-day level. We define treated "early-month" days as days 1-3 of each month. These coincide with recurring schedules, including Supplemental Security Income payments on the 1st, some payroll dates and rent due on the 1st. The design does not identify a single mechanism. Days 10-24 are not free of payment dates (Social Security's Wednesday payments, among others, fall within them), and the contrast estimates a month-boundary excess rather than an effect of payment. January was excluded because the New Year's Eve surge distorts the ordinary early-month pattern. Control days were same-day-of-week dates on days 10-24 of the same month-year, which lie farthest from both month

boundaries (calls stayed elevated for about a week after the 1st and ran low late in the preceding month; Fig. 2b) and include every weekday at least twice. We estimated conditional Poisson models[36,37] with matched-set fixed effects defined by month-year x day-of-week (132 month-year cells producing 396 case-day matched sets across the 2012-2023 panel, with January excluded), and included temperature as a covariate. The primary estimand was the incidence rate ratio for days 1-3 versus matched later-month controls. We applied the same design to comparison endpoints and all dispatcher-coded calls. To model category share, we fitted a binomial fixed-effects model to the drug or poisoning count versus all other calls, using the all-call denominator as the binomial weight and the same treatment, temperature, matched-stratum fixed effects and clustered inference. A full-raw-file construct audit varied the drug or poisoning code definition (Extended Data Table 3; Supplementary Methods). Incident-level disposition and acuity models used the same matched strata, ZIP-hour weather and date-clustered inference. Secondary checks used the same matched strata: mid-month placebos, four-year subperiods, event-time curves by borough and ZIP-income quartile compared with 99,999 circular-shift permutations, six-hour blocks with hour-specific matched sets, Social Security Wednesday and payment-advance contrasts, and call-type transitions, dispositions and endpoint-specific neighbourhood slopes (Extended Data Table 4; Supplementary Fig. S2; Supplementary Tables S7, S10, S15, S30 and S44).

Holiday event windows

We analysed holiday windows using matched hourly comparisons. For each holiday occurrence and each relative hour in the event window, we matched the same relative hour 7 and 14 days before and after the holiday. Candidate control anchors or matched dates that overlapped other major holidays were excluded. July 4 was analysed using windows for the full holiday day, daytime, evening, the overnight hours after midnight, and a prespecified fireworks window spanning 18:00 on July 4 through 05:59 on July 5. Memorial Day and Labor Day were analysed as a second holiday class using full-day, overnight, daytime, afternoon, and evening windows. Window summaries were estimated with conditional Poisson models including matched-set fixed effects, temperature, and precipitation covariates, with standard errors clustered by holiday occurrence. Because the cluster count is small (G = 12 for July 4, G = 24 for pooled Memorial and Labor Day), we re-estimated the eight primary endpoint x event cells using a paired event-cluster bootstrap (B = 999) and a t critical value with G - 1 degrees of freedom as a small-cluster sensitivity (Extended Data Table 5). For those eight cells, binomial share models contrasted each mutually exclusive final-code endpoint with all other calls using the all-call denominator, the same treatment, weather covariates, matched-set fixed effects and event clusters, with the same t intervals. Exploratory endpoint-by-window tests used Benjamini-Hochberg adjustment. These analyses were not preregistered; the primary endpoints (burns, assault, drug or poisoning, with cardiac arrest as a specificity null) and fireworks window were fixed before the exploratory sweep. All other tests and their adjusted values are reported in Supplementary Table S2.

Protocol-locked calendar validation in 2024-2025

Before extracting endpoint-specific 2025 counts and applying the cached 2024-2025 response-time extract, we fixed outcomes, windows and decision rules in a dated protocol. Because 2024 endpoint counts had appeared in an exploratory update, this exercise was only partly prospective; the 2025 counts were untouched. No later result altered endpoints or windows. After the directional result was known, a no-intercept calibration slope of 2024-2025 on 2012-2023 log-IRRs summarized magnitude (Supplementary Table S39). A design fixed before fitting evaluated,

without refitting, how much 2024-2025 within-month daily load nested 2012-2023 models anticipated from the weekday-by-month rhythm, calendar pulse terms and realized weather (Supplementary Table S43).

We aggregated the same dispatcher codes for drug or poisoning, assault, mental health, cardiac arrest and burns to city-hour counts. The early-month contrast compared days 1-3 with same-weekday days 10-24 within each non-January month-year. July 4 used the 18:00-05:59 fireworks window and same-hour controls at +/-7 and 14 days. Memorial and Labor Day used the full holiday day with the same shifted controls. Both periods used the same unadjusted matched contrast, so the direction test did not depend on weather series of different vintages; the 2024-2025 estimates are directional validation contrasts, not weather-adjusted magnitudes.

The seven locked nonzero directions were early-month drug or poisoning positive; July 4 burns, assault and drug or poisoning positive; and Memorial/Labor assault and drug or poisoning positive with mental health negative. The count pattern was supported if at least six directions matched, including early-month drug or poisoning, with no holiday class reversing both assault and drug or poisoning. The locked calendar-event response-time classifier retained total-call-weighted city-hour means, required an early-month estimate no larger than +30 seconds and negative estimates for both holiday classes, and failed after the July 4 reversal. With only two July 4 and four Memorial/Labor occurrences, these were directional criteria. Full rules, secondary summaries and the separate permit condition are reported in Supplementary Methods and Table S18.

Frozen external-city calendar validation

We tested spatial recurrence in three public incident-level Fire/EMS sources. City selection reflected the availability of public incident records with usable timestamps and mappable call categories. Phoenix contributed 2014-2018 calls classified by the source as advanced or basic life support. Glendale contributed 2021-2025 incidents within the City of Glendale whose populated call types mapped to medical categories in the shared 2017-2018 Phoenix regional code dictionary; a source outage (23-29 October 2023, recorded as zero calls) was excluded by documented amendment. Cincinnati contributed all 979,415 Fire/EMS CAD incidents in complete years 2015-2025. Source timestamps were interpreted as local wall time (America/Phoenix or America/New_York). The analysis code verified row identity, annual call-type coverage and frozen file hashes before constructing exposures. Endpoint mappings were fixed from each dispatch vocabulary before estimation. The Cincinnati crosswalk was required to have at least 0.98 sensitivity and positive predictive value against populated official categories; the observed minima were 0.996 and 0.990. No defensible common mental-health category was available, so that endpoint was not transported. The full mappings and exclusions are reported in Supplementary Methods.

The early-month, July 4 and Memorial/Labor windows and shifted controls were identical to the locked New York City count design. We fitted city-specific conditional Poisson models with matched-stratum fixed effects. Early-month standard errors were clustered by month-year-weekday stratum. Holiday standard errors were clustered by event occurrence, and intervals used a t reference with G - 1 degrees of freedom. The primary models were unadjusted to retain a common specification across weather vintages. Direction, not statistical significance, was the frozen replication criterion. We report every target and cardiac-arrest comparison cell and do not pool cross-city magnitudes (Supplementary Table S24).

Marathon route-disruption design

The marathon analysis used a difference-in-differences design[38] with incident-level EMS response times on held NYC Marathon Sundays in 2013-2019 and 2021-2023, together with matched control Sundays 2 and 4 weeks before and after each held event. The course was reconstructed along its documented streets by routing between labelled waypoints on OpenStreetMap; it measured 42.5 km against the 42.2-km course, within a 2% length tolerance fixed in advance (Supplementary Methods). ZIP codes were classified by the distance from their 2020 ZIP Code Tabulation Area polygon centroids to the course: affected ZIPs had centroids within 1 km, control ZIPs were more than 5 km away, and ZIPs between 1 and 5 km were excluded from the route contrast. The primary access outcome was call-to-scene time (dispatch plus travel), available for 94,768 of 95,996 candidate incidents (98.7%) and not conditional on completed transport. The secondary complete call-to-hospital outcome comprised 65,639 incidents across 118 ZIP codes, 10 marathon Sundays and 40 matched control Sundays. Records missing any interval component or with non-positive or extreme totals were excluded from that secondary outcome (31.6% of candidates, chiefly incidents without a completed transport), and sensitivity to extreme values is reported in Supplementary Table S6. The complete interval was analysed first; prioritizing call-to-scene was not preregistered. A dated phase specification was then frozen before fitting: prior Saturday 21:00-23:59, Sunday setup 00:00-06:59, active operations 07:00-20:59 and post-event 21:00-23:59. Some closures precede 07:00, so early Sunday is not a clean pre-closure placebo (Extended Data Fig. 1; Supplementary Methods).

We modelled differential interval validity with the design fixed effects. Lower call-to-scene validity near the route on marathon days prompted post hoc selection checks. In the full-day and equal-weight active-window models, we varied the upper validity ceiling and assigned invalid intervals over an adverse 1-3,600-second grid. These deterministic stress tests are not multiple imputation; their intervals condition on the imposed values (Supplementary Tables S19 and S23). Further sensitivities excluded records flagged by the source as standby and adjusted for initial call type and severity (+161 seconds, 95% CI 64-258); the latter conditions on dispatch-visible case mix that the event could itself change and is not the preferred causal model (Supplementary Table S23). A check designed after the main estimate was known repeated the equal-year estimate for initial call types unlikely to arise from runners or spectators, for call types whose route-adjacent counts did not rise on race days and for life-threatening calls (initial segments 1-3), excluding special-event and standby records (Supplementary Table S48).

The all-day model used ZIP, year-month and hour fixed effects, a temperature spline, precipitation and ZIP-and-date clustering. The principal phase model instead used ZIP by hour-of-week and physical UTC-hour fixed effects. We summarized ten annual effects equally using t(9) inference, leave-one-event omission and 19,999 matched-date permutations (Supplementary Table S23). Every held race fell on the Sunday daylight saving time ended. Active and prior-Saturday windows avoid the ambiguous fall-back hour, and 1 November 2020, a fall-back Sunday on which the race was virtual and no streets closed, served as a clock-change placebo analysed with the same event-effect arithmetic (Supplementary Table S29). A check fixed before estimation used only standard-time or only daylight-time comparison Sundays and treated the 12 spring-forward Sundays of 2012-2023 as placebos (Supplementary Table S49). Supplementary Methods detail daylight-saving sensitivities and covariance diagnostics. Fire-company responses on the same Sundays (Fire Incident Dispatch Data) were analysed with the same design, separately for non-medical incidents and medical co-responses (Supplementary Table S27). A frozen port of the design to Cincinnati's Flying Pig Marathon used official closure geometry for 2022-2025; its extension to six earlier races was

conditioned on a prespecified selection gate evaluated before any response time was computed (Supplementary Table S42).

Secondary analyses evaluated complete-interval components, income, route distance, the cancelled 2012 event, quantiles[39] and response thresholds (Extended Data Fig. 2 and Tables 7 and 8; Supplementary Tables S6, S8 and S9).

United Nations General Assembly weeks

We applied a complementary difference-in-differences design to United Nations General Assembly high-level weeks. Exposure comprised the first five weekdays of the General Debate each year (official opening dates, 2012-2023); control days were the same weekdays 2 and 4 weeks before and after. Treated ZIPs were Manhattan ZIPs with centroids within 1.5 km of UN headquarters (10016, 10017, 10022); ZIPs 1.5-3 km away were excluded as a buffer and the remaining Manhattan ZIPs served as controls. The primary outcome was ZIP-hour mean call-to-scene time among positive intervals below 7,200 seconds, weighted by the number of valid intervals. Dispatch, travel, scene and transport intervals were modelled separately. Models included ZIP and event-year fixed effects, hour and day-of-week adjustment, and standard errors clustered by ZIP and date. We also estimated each in-person event year separately and summarized the 11 annual contrasts equally with t inference using 10 degrees of freedom. With only three treated ZIPs, within-year intervals are unreliable, and the equal-year summary treats years as the units of inference. The 2020 General Debate, held virtually without street closures, provided a placebo year (Supplementary Table S11); the 2021 session, held in hybrid form with leaders attending in person, was classed as in-person. We applied the same fixed-effects structure to the fraction of calls with a valid call-to-scene interval as a selection diagnostic (Supplementary Table S23).

Permitted street-event census

NYC's historical permitted-event records provided a separate generalization test. The permit census defined all 2024-2025 same-day events before examining EMS outcomes. The primary permit class comprised full-street closures for parades or athletic races spanning at least two police precincts; local full closures were single-precinct block parties or single-block festivals, and less-restrictive permits formed a third class. Marathon, UN General Assembly and Open Streets records were excluded by name. We modelled valid-response-count-weighted precinct-hour means with precinct-by-hour-of-week, precinct-by-year-month and borough-by-date-by-hour fixed effects, clustering by precinct and date. The primary sample excluded precincts 105, 113 and 116 because the 116th Precinct opened during the 2024-2025 observation period; standby records and source-flagged special-event assignments were excluded from outcome aggregation. Prespecified checks modelled call-to-scene validity and the dispatch, assignment-to-scene travel and scene intervals separately; shifted exposure to the preceding window and +/-7 and 14 days; retained the precinct split; and repeated estimates by year and with each event date omitted (Supplementary Table S16). This design estimates the association with the permit-defined event class, not the isolated effect of physical closure. A design fixed before fitting added moderators to the same model: expected citywide demand for each hour, predicted from hours without multi-precinct events, and holiday weekends; a precinct crosswalk distributed the pooled excess for a descriptive equity summary (Supplementary Table S45). A further design fixed before fitting repeated the model without federal-holiday event dates, without Manhattan

precincts, without both and by year, with a prespecified generalization rule for the estimate without both (Supplementary Table S49).

After the parent result, an exploratory scale extension restricted to full-street parades and athletic races compared exposure spanning one, two or at least three precincts (Supplementary Table S17).

Joint changes in demand and access

Complementary models estimated marathon call volume and calendar-event response time. Demand-stratified models tested whether the marathon access penalty was larger during higher baseline demand (Supplementary Tables S12 and S13). A design fixed after these results were known multiplied incidents by mean valid call-to-scene time within each schedule's reported comparison; incidents without a valid interval take the cell's mean wait, and the excess divides exactly into parts from more calls, slower responses and their interaction (Supplementary Methods; Supplementary Table S46). Valid call-to-scene time is dispatch plus travel, more than 0 and under 3,600 seconds. For an exposed cell (1) and its counterfactual (0), with N incidents at mean wait T, the excess is (N1 - N0)T0 + N0(T1 - T0) + (N1 - N0)(T1 - T0): more calls at the usual wait, the usual calls at the changed wait, and their interaction. Pooled schedules average these parts over occurrences, so the interaction also absorbs covariance across occurrences. Street operations use a near-far counterfactual: the wait adds the distant-ZIP change to the route-adjacent control mean, and the count scales the route-adjacent control count by the distant-ZIP ratio, so the wait difference equals the reported event effect. The relative change divides the excess by mean comparison waiting; intervals are percentiles of 2,000 cluster bootstrap draws over month-years, occurrences, races or years.

Independent transport corroboration

After the parent EMS results were known, we fixed an exploratory extension before inspecting its new outcomes. Official high-volume for-hire vehicle records from all 24 months of 2024-2025 supplied pickup time, endpoint taxi zones, duration and distance[40]. We compared identical origin-destination and pickup-hour cells on event dates and uncontaminated same-hour dates +/-7 and 14 days. Fixed control-trip weights standardized the retained cells before exposed-minus-unexposed differencing. The equal-event summary used 1,999 date-cluster bootstrap draws. An analogous EMS contrast used baseline valid-response weights; its correlation with trip duration is descriptive. Zone assignment, local civil clocks, daylight-saving exclusions, support rules and retention are specified in Supplementary Table S26. A comparison-trained sensitivity paired event/placebo cells and weights (Supplementary Table S28).

Joint calendar and neighbourhood models

These designs were fixed and hashed before fitting. A conditional quasi-Poisson model of city-day counts for 2012-2023, January included, estimated all calendar pulses together, with year-month strata, day-of-week effects specific to each calendar month, day-of-month terms for days 1-9 and 25-31 against days 10-24, indicators from 3 days before to 7 days after New Year's Day, July 4, Memorial Day and Labor Day, distributed-lag temperature (lags 0-3)[41], precipitation and a day-length offset, with year-month clustering; each holiday was compared with the same month's Sunday-versus-weekday coefficient. ZIP-specific pulse estimates were pooled in a two-stage multivariate meta-analysis[42] and related to income, race and ethnicity (American Community Survey 2015-2019), public housing, density and licence densities, with Benjamini-Hochberg control. Supplementary Tables S31-S34 give full

specifications. Further designs fixed before fitting covered overlap tests, a holiday day-off decomposition, a marathon held-call check, borough and benefit-receipt sensitivities, a weekly Memorial and Labor Day dose-response check, a population-presence check with Advan foot traffic (2019-2023) and a parade check excluding six ZIPs along the Labor Day route (Supplementary Methods; Supplementary Tables S31, S33, S36-S38 and S40).

Two further designs, fixed and hashed on 11 November 2026 before fitting, tested whether each pulse's ZIP map reproduced itself between odd and even years and whether the pulses shared one neighbourhood profile (Supplementary Table S50). Bivariate random-effects models compared the maps, a multivariate meta-regression compared slopes scaled by each pulse's between-ZIP standard deviation, and an incident-level model added the neighbourhood measures and hospital access to the marathon design.

The presence check measured daily Advan Weekly Patterns Plus visits per point of interest in the 170 panel ZIPs (2019-2023) and refitted the early-month and holiday call models without and with log presence (Supplementary Table S38).

Statistical software and inference

Analyses used R (data.table, fixest, dlnm, mixmeta, spdep) and Python (pandas, numpy, duckdb, geopandas). Coefficient and matched-date permutation tests were two-sided; circular-shift shape and Moran's I permutation tests were one-sided. Joint Wald tests used the upper chi-squared or F tail. Confidence intervals used design-specific clustered variance estimators unless noted.

Use of generative artificial intelligence

Claude Code (Anthropic) and Codex (OpenAI), accessed from June through November 2026, were used to assist with analysis code, audit scripts, figure code and manuscript editing. The authors reviewed generated code against authoritative tabular outputs, reran the reported analyses and take full responsibility for the content.

Data availability

All analyses used publicly available data except the population-presence check (Supplementary Table S38), which used Advan Weekly Patterns Plus foot-traffic data licensed through Dewey. Restrictions apply to the availability of these data, which were used under licence for this study and cannot be redistributed; they are available from the provider under its terms.

The NYC EMS Incident Dispatch Data are publicly available through NYC Open Data under dataset ID 76xm-jjuj. The primary analyses use the frozen 14 June 2024 extract EMS_Incident_Dispatch_Data_20240614.csv (25,984,643 rows excluding the header; SHA-256 f70928cbacd8b33b0bc066e8af392402875cfd0878fa8677d83d54ba0e47a271). The extraction date is known; the Socrata rowsUpdatedAt value at extraction was not captured and is not inferred from the live portal. Hourly weather data came from the North American Land Data Assimilation System (NLDAS-2 forcing, product NLDAS_FORA0125_H.002, NASA Goddard Earth Sciences Data and Information Services Center) and were linked by ZIP code and UTC hour. ZIP income strata used American Community Survey 2015-2019 estimates for ZIP Code Tabulation Areas with land in the five NYC counties; distance calculations used the 2020 Census Gazetteer ZIP Code Tabulation Area file. Marathon route distances were computed from the course reconstructed along its documented streets from OpenStreetMap (© OpenStreetMap contributors, Open Database License) with

FOSSGIS OSRM routing; the geometry is included in the archive. FDNY Fire Incident Dispatch Data (NYC Open Data 8m42-w767) were retrieved on 19 October 2026 for the marathon and matched control dates. Separate 2024-2025 EMS extracts cover 1 January 2024 through 31 December 2025 and were frozen on 13 and 18 July 2026; the 2024-2025 marathon and UN incident extracts were retrieved in early October and on 2 November 2026. The calendar-pattern validation uses the city-hour endpoint and response-time extracts; the permit-census generalization test uses the precinct-hour response and count extracts. Permit records came from the NYC Permitted Event Information dataset bkfu-528j; its 21,748 exact-distinct 2024-2025 rows were frozen on 13 July 2026 before EMS outcomes were examined. Police precinct boundaries came from NYC Open Data y76i-bdw7. File sizes, row counts, MD5 hashes, query identifiers and exclusion rules are reported in Supplementary Table S22 and the versioned release manifest.

The external calendar sources were the City of Phoenix Fire calls-for-service files (https://www.phoenixopendata.com/dataset/caf49f72-f22f-4ad9-9405-2a3db9619423), the City of Glendale historic fire-incident layer (https://services1.arcgis.com/9fVTQQSiODPjLUTa/arcgis/rest/services/Glendale_Fire_Incidents_Historic/FeatureServer/0) and the City of Cincinnati Fire/EMS CAD dataset vnsz-a3wp (https://data.cincinnati-oh.gov/d/vnsz-a3wp). Cincinnati's 2015-2025 snapshot was frozen on 14 August 2026; source hashes, row accounting and endpoint crosswalk checks are reported with Supplementary Table S24. Flying Pig Marathon closures came from City of Cincinnati Department of Transportation and Engineering closure notices and Cincinnati Area Geographic Information System (CAGIS) street and boundary services.

The transport extension used the New York City Taxi and Limousine Commission (TLC) monthly high-volume for-hire vehicle Parquet files for January 2024-December 2025 and published taxi-zone geometry[40]. Retrieval dates, exact queries, source hashes and event-support ledgers accompany Supplementary Table S26.

The joint neighbourhood models used American Community Survey 2015-2019 race and ethnicity counts and household benefit receipt (tables B03002, B19056, B19057 and B22010) for ZIP Code Tabulation Areas (ZCTAs), New York State Liquor Authority active licences (9s3h-dpkz), New York State Department of Health hospital and emergency-department certification records (vn5v-hh5r, 2g9y-7kqm), New York City Housing Authority (NYCHA) development addresses and data book (3ub5-4ph8, evjd-dqpz), Department of Consumer and Worker Protection pawnbroker licences (w7w3-xahh) and NYC Department of Health and Mental Hygiene (DOHMH) modified ZCTA geometry. These sources were retrieved on 15 and 19 October 2026; file hashes accompany the release.

Code availability

Code, result tables, figures, dated designs, quantitative claim traces, source contracts and verification commands accompany this manuscript for peer review as calendar_rhythms_v3.4.6_code_results_review.zip. This reduced review distribution of study 3.4.6 omits bulk input data and cached data/model objects. Source identifiers, retrieval scripts, recorded data vintages and restoration paths are documented. A current NYC Open Data download may differ from the frozen input vintage; the original incident-level EMS snapshot requires archive restoration for a full raw-data rerun. The full study archive, including frozen aggregate inputs, and its separately pinned optional MCP companion remain available to editors and reviewers on request. Permanent DOI deposition is pending.

The joint calendar and neighbourhood models and the October and November analyses (Supplementary Tables S27, S29 and S31-S50) are included with their frozen designs, source/retrieval records and outputs. The licensed Advan input to Supplementary Table S38 is not redistributed, but its code, design and estimates are included. The source contract documents the restoration requirements for the external-city inputs.

A self-contained interactive visualization, its source tables and verification tools are included in the review distribution.

Ethics statement

This study used publicly available, de-identified data and licensed aggregate foot-traffic counts, and was determined exempt by the University of Arizona Institutional Review Board.

Funding

This work was supported by the National Institute of Environmental Health Sciences (R01ES034740).

Author contributions

C.C.L. conceptualized the study, designed the analysis and wrote the manuscript. R.W. contributed to spatial analysis methods and the marathon design. J.A.R. contributed to study design and data interpretation. H.K. contributed to study design and statistical methodology. X.L. contributed to environmental exposure assessment. L.G. contributed to epidemiological design and data interpretation. C.C. contributed to study design and data interpretation. N.C.M. contributed to study design, provided EMS domain expertise and contributed to data interpretation. All authors reviewed and approved the final manuscript.

Competing interests

The authors declare no conflicts of interest.

References

1. Reades, J., Calabrese, F., Sevtsuk, A. & Ratti, C. Cellular census: Explorations in urban data collection. *IEEE Pervasive Computing* **6**, 30–38 (2007).

2. Calabrese, F., Colonna, M., Lovisolo, P., Parata, D. & Ratti, C. Real-time urban monitoring using cell phones: A case study in Rome. *IEEE Transactions on Intelligent Transportation Systems* **12**, 141–151 (2011).

3. Miranda, F. *et al.* Urban pulse: Capturing the rhythm of cities. *IEEE Transactions on Visualization and Computer Graphics* **23**, 791–800 (2017).

4. Wang, L., Qian, C., Kats, P., Kontokosta, C. & Sobolevsky, S. Structure of 311 service requests as a signature of urban location. *PLOS ONE* **12**, e0186314 (2017).

5. Sirenko, M., Comes, T. & Verbraeck, A. The rhythm of risk: Exploring spatio-temporal patterns of urban vulnerability with ambulance calls data. *Environment and Planning B: Urban Analytics and City Science* **52**, 863–881 (2025).

6. Noulas, A., Moffatt, C., Hristova, D. & Gonçalves, B. Foursquare to the rescue: Predicting ambulance calls across geographies. in *Proceedings of the 2018 international conference on digital health* 100–109 (ACM, 2018). doi:10.1145/3194658.3194665.

7. Lowthian, J. A. *et al.* Increasing utilisation of emergency ambulances. *Australian Health Review* **35**, 63–69 (2011).

8. Mann, N. C. & Hedges, J. R. The role of prehospital care providers in the advancement of public health. *Prehospital Emergency Care* **6**, S63–S67 (2002).

9. Salley, C., Mohammadi, N., Taylor, J. E. & Bozeman, J. F., III. Inequities in urban emergency medical services. *npj Health Systems* **3**, 92 (2026).

10. Chikish, Y. *et al. Congestion Pricing and Emergency Medical Service Response: Evidence from New York City*. (2026) doi:10.3386/w35414.

11. Besedovsky, N., Grafe, F.-J., Hilbrandt, H. & Langguth, H. Time as infrastructure: For an analysis of contemporary urbanization. *City* **23**, 580–588 (2019).

12. Zerubavel, E. *Hidden Rhythms: Schedules and Calendars in Social Life*. (University of Chicago Press, Chicago, 1981).

13. Sharma, S. *In the Meantime: Temporality and Cultural Politics*. (Duke University Press, Durham, NC, 2014).

14. Phillips, D. P., Christenfeld, N. & Ryan, N. M. An increase in the number of deaths in the United States in the first week of the month: An association with substance abuse and other causes of death. *New England Journal of Medicine* **341**, 93–98 (1999).

15. Dobkin, C. & Puller, S. L. The effects of government transfers on monthly cycles in drug abuse, hospitalization and mortality. *Journal of Public Economics* **91**, 2137–2157 (2007).

16. Stephens, M., Jr. '3rd of tha month': Do social security recipients smooth consumption between checks? *American Economic Review* **93**, 406–422 (2003).

17. Riddell, C. & Riddell, R. Welfare checks, drug consumption, and health: Evidence from Vancouver injection drug users. *Journal of Human Resources* **41**, 138–161 (2006).

18. Foley, C. F. Welfare payments and crime. *Review of Economics and Statistics* **93**, 97–112 (2011).

19. Evans, W. N. & Moore, T. J. The short-term mortality consequences of income receipt. *Journal of Public Economics* **95**, 1410–1424 (2011).

20. Brunette, D. D., Kominsky, J. & Ruiz, E. Correlation of emergency health care use, 911 volume, and jail activity with welfare check distribution. *Annals of Emergency Medicine* **20**, 739–742 (1991).

21. Verheul, G., Singer, S. M. & Christenson, J. M. Mortality and morbidity associated with the distribution of monthly welfare payments. *Academic Emergency Medicine* **4**, 118–123 (1997).

22. Arbon, P. The development of conceptual models for mass-gathering health. *Prehospital and Disaster Medicine* **19**, 208–212 (2004).

23. Jena, A. B., Mann, N. C., Wedlund, L. N. & Olenski, A. Delays in emergency care and mortality during major U.S. marathons. *New England Journal of Medicine* **376**, 1441–1450 (2017).

24. Kienbacher, C. L. *et al.* Social mass gathering events influence emergency medical services call volume. *Frontiers in Public Health* **12**, 1394384 (2024).

25. Apiratwarakul, K., Cheung, L. W., Gaysonsiri, D. & Ienghong, K. Recognizing the demand for emergency medical services during the Songkran festival, Thailand's intangible cultural heritage: An approach for healthcare policy development. *Critical Public Health* **35**, 2529379 (2025).

26. Blackburn, C. C. & Rico, M. Impact of collegiate football games on emergency response intervals: A case study of College Station, Texas (USA). *Prehospital and Disaster Medicine* **39**, 270–274 (2024).

27. City of New York. EMS incident dispatch data (NYC Open Data dataset 76xm-jjuj). (2024).

28. Batty, M. *The New Science of Cities*. (MIT Press, Cambridge, MA, 2013).

29. Pons, P. T. & Markovchick, V. J. Eight minutes or less: Does the ambulance response time guideline impact trauma patient outcome? *Journal of Emergency Medicine* **23**, 43–48 (2002).

30. Nicholl, J., West, J., Goodacre, S. & Turner, J. The relationship between distance to hospital and patient mortality in emergencies: An observational study. *Emergency Medicine Journal* **24**, 665–668 (2007).

31. U.S. Consumer Product Safety Commission. 2019 fireworks annual report: Fireworks-related deaths, emergency department-treated injuries, and enforcement activities during 2019. (2020).

32. Maassel, N. *et al.* Firework-related injuries treated at emergency departments in the United States during the COVID-19 pandemic in 2020 compared to 2018–2019. *Injury Epidemiology* **8**, 65 (2021).

33. Phillips, D. P., Jarvinen, J. R., Abramson, I. S. & Phillips, R. R. Cardiac mortality is higher around Christmas and New Year's than at any other time: The holidays as a risk factor for death. *Circulation* **110**, 3781–3788 (2004).

34. Landrigan, D. FDNY orders ambulances to transport patients to nearest ER, drawing criticism. *JEMS* (2025).

35. Wheeler, E., Jones, T. S., Gilbert, M. K. & Davidson, P. J. Opioid overdose prevention programs providing naloxone to laypersons in the United States, 2014. *MMWR. Morbidity and Mortality Weekly Report* **64**, 631–635 (2015).

36. Maclure, M. The case-crossover design: A method for studying transient effects on the risk of acute events. *American Journal of Epidemiology* **133**, 144–153 (1991).

37. Janes, H., Sheppard, L. & Lumley, T. Case-crossover analyses of air pollution exposure data: Referent selection strategies and their implications for bias. *Epidemiology* **16**, 717–726 (2005).

38. Angrist, J. D. & Pischke, J.-S. *Mostly Harmless Econometrics: An Empiricist's Companion*. (Princeton University Press, Princeton, NJ, 2009).

39. Firpo, S., Fortin, N. M. & Lemieux, T. Unconditional quantile regressions. *Econometrica* **77**, 953–973 (2009).

40. New York City Taxi and Limousine Commission. TLC Trip Record Data. (2026).

41. Gasparrini, A., Armstrong, B. & Kenward, M. G. Distributed lag non-linear models. *Statistics in Medicine* **29**, 2224–2234 (2010).

42. Sera, F., Armstrong, B., Blangiardo, M. & Gasparrini, A. An extended mixed-effects framework for meta-analysis. *Statistics in Medicine* **38**, 5429–5444 (2019).

Figure Legends

Figure 1. Pre-arrival EMS access during street-event operations. (a) The marathon course, reconstructed along its documented streets, with ZIP codes shaded by distance from it; legend values are the active-window own changes in mean call-to-scene time for each distance band, from the ZIP-hour means in b. Borough outlines are drawn from ZIP boundaries and are approximate; ZIPs 11249 and 11430 have no separate polygon in the city's modified ZIP Code Tabulation Areas and are not drawn. (b) Own change in call-to-scene time on race Sundays relative to their matched Sundays, by distance from the course, from ZIP-hour means weighted by total calls (triangles) and from the headline incident sample for the route-adjacent and distant arms (circles). The ZIP-hour bands use the 170-ZIP weather-linked panel (24 ZIPs within 1 km), whereas the incident arms use every ZIP (34 within 1 km). The ZIP-hour estimator has no validity filter, so its magnitudes exceed the incident-level estimates; the headline +152 s is the difference between the two incident-level estimates. Equal-event means with t intervals on 9 degrees of freedom. (c) Call-to-scene estimates on a common scale, with 95% confidence intervals. The marathon active-window summary gives equal weight to 10 annual contrasts and uses t inference with 9 degrees of freedom (+152 s, 95% CI 93-210; P = 0.00024); the dispatch and travel rows are equal-year estimates for those intervals in the principal sample and sum to the call-to-scene estimate up to rounding. The prior-Saturday placebo (grey) uses the incident-level fixed-effects specification, in which the active-minus-prior-Saturday difference is +138 s (95% CI 54-221; Supplementary Table S23). The United Nations General Assembly (UNGA) summary gives equal weight to 11 in-person years (t inference with 10 degrees of freedom, P = 0.00012). The permit rows cover permitted street events in 2024-2025, classified before EMS outcomes were examined, and test whether access findings extended to other scheduled street events; they use 532,122 precinct-hours and precinct-and-date-clustered intervals (multi-precinct full-street parades/races: P = 0.0011). Permit estimates describe an event class; they do not identify an effect of spatial scale or isolate physical closure from other operations. Filled symbols mark intervals that exclude 0 and open symbols intervals that include 0. Annual estimates are in Supplementary Fig. S5a.

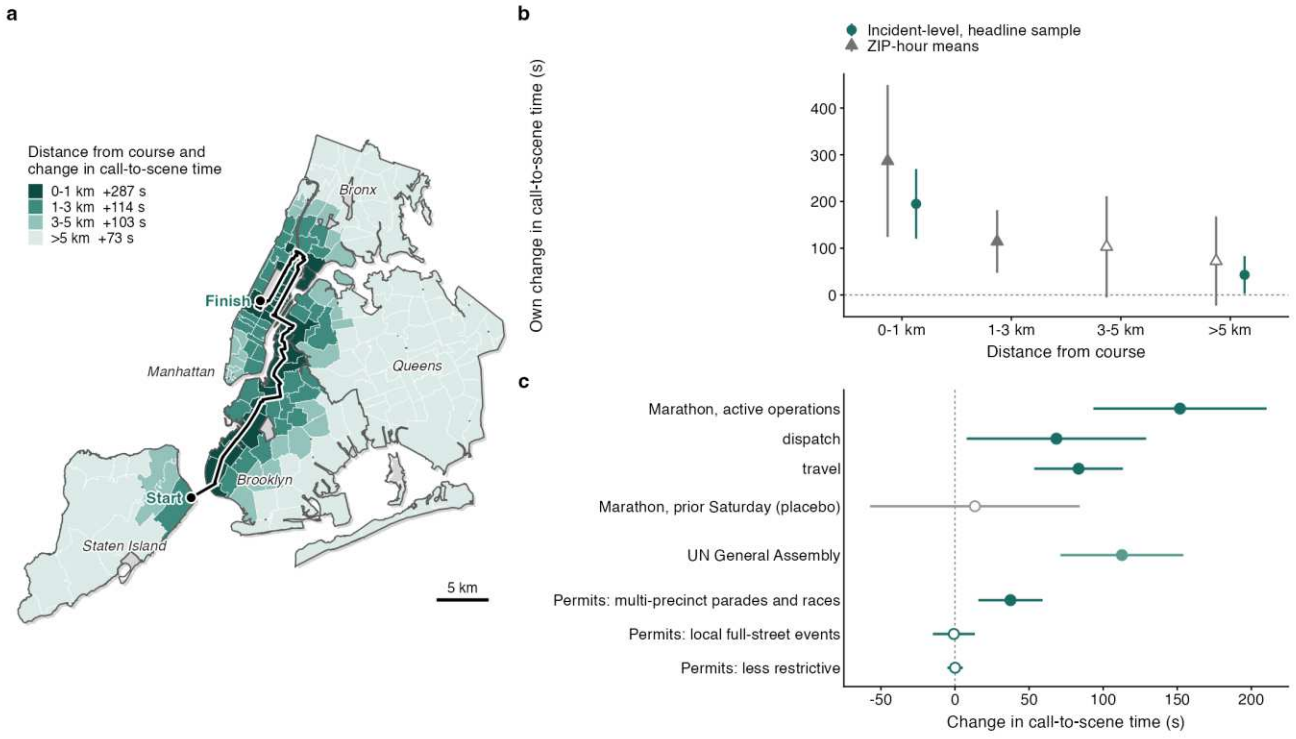


Figure 2. Where and when the early-month pulse in drug or poisoning calls occurred. (a) Increase in drug or poisoning calls on days 1-3 against days 10-24 in each ZIP code, February-December 2012-2023, shown as best linear unbiased predictions from the two-stage multivariate meta-analysis, which shrink each ZIP's estimate towards the citywide mean in proportion to its uncertainty; dots mark the 58 ZIPs whose own 95% interval excludes 0. Estimates attach to incident location and are ecological. Borough outlines are approximate; ZIPs 11249 and 11430 are not drawn. (b) Model-based event-time incidence rate ratios from 4,011 city-days, with days 10-24 as the reference period; vertical lines show 95% confidence intervals, filled points mark days whose interval excludes 1, and the shaded region marks days 1-3; negative days count back from the 1st. (c) Day-of-month terms by ZIP income quartile from the stacked model, against days 10-24; the day 1-9 curves differ across quartiles ($P < 0.001$) and the day 25-31 terms do not ($P = 0.50$). Filled points mark 95% confidence intervals that exclude 0. Dates use America/New_York local time. Descriptive daily means and the pooled count, share and 2024-2025 estimates are in Supplementary Fig. S5b,c.

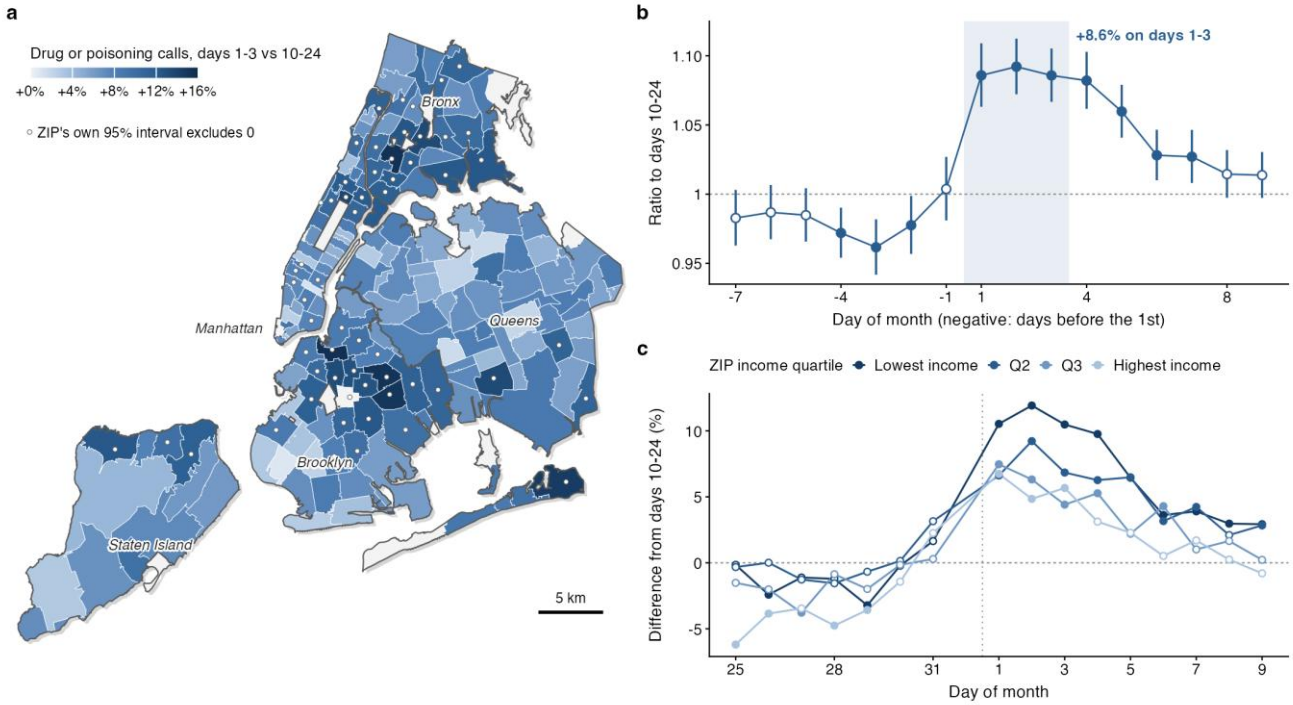


Figure 3. Holiday calls hour by hour. (a) Hourly rate ratios for seven call categories from noon on July 4 to 05:59 on July 5, each against the same hour 7 and 14 days before and after, pooled over 12 occurrences (2012-2023); dots mark hours whose 95% interval excludes 1, and the bar marks the prespecified fireworks window. Hourly estimates are descriptive and not adjusted for multiple comparisons. (b) Call-count incidence rate ratios (circles) and category-share odds ratios (squares) during 12 July 4 fireworks windows and 24 pooled Memorial/Labor Day occurrences. Share estimates compare the adjusted odds of each category versus all other calls between event and control windows; they are not percentage-point changes in share. Lines show 95% confidence intervals based on event-clustered models and t critical values with 11 or 23 degrees of freedom; filled symbols mark intervals that exclude 1. A category's share can change while its count does not, because the share's denominator includes the other calls that rose or fell; this is why cardiac-arrest shares moved without a detected change in counts.

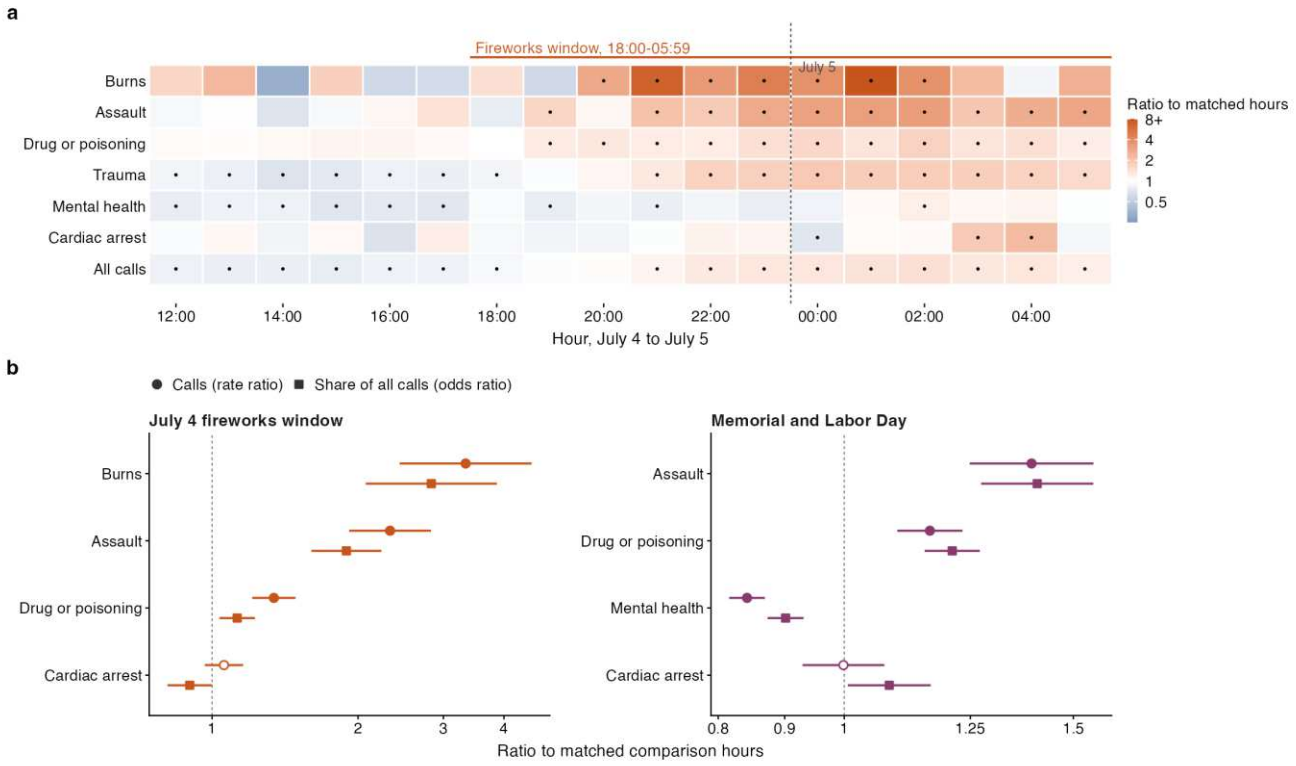


Figure 4. Each schedule's geography. (a) ZIP codes within 1 km of the reconstructed marathon course and beyond 5 km from it, the two arms of the marathon incident model; ZIP codes 1-5 km away are not in the contrast. (b-d) Each ZIP code's pulse in drug or poisoning calls for the early month (b), Memorial and Labor Day (c) and July 4 (d), relative to the pooled estimate across 170 ZIP codes and expressed in standard deviations of that pulse's variation between ZIP codes. Values are best linear unbiased predictions from the two-stage multivariate meta-analysis (Supplementary Table S32); values beyond ±2 are shown at the limit, and the thin line marks the marathon course. Labels give the correlation of each pulse's true ZIP effects between odd and even years; the July 4 interval included 0, so its map was too unstable for the prespecified comparison with the other pulses' maps. Borough outlines are approximate; ZIPs 11249 and 11430 are not drawn. (e) What each pattern followed. Marathon delay (left): added call-to-scene time in route-adjacent ZIP codes (course row) and per standard deviation of each neighbourhood measure, from an incident-weighted model that also holds distance to the nearest emergency department, and whether the line to it crosses the course, fixed, with t intervals on 9 degrees of freedom. The shaded band marks the prespecified bound of ±50 s per standard deviation for income, racial composition and public housing. Calendar pulses (right): change per standard deviation of each measure divided by the pulse's standard deviation between ZIP codes, with percentile intervals from 499 ZIP bootstrap resamples; on the course row the sign of the distance slope is reversed, so positive values mean a larger pulse nearer the course. Labels give the share of each pulse's between-ZIP variance explained by income, racial composition and public housing together, as point estimates. Because the July 4 map is unstable, sampling noise biases its slopes and explained share toward zero. Measures are log median income, non-Hispanic Black share, share of residents in public housing and log population density, each standardized across ZIP codes; the models also hold pawnbroker licence density fixed. Filled symbols mark 95% intervals that exclude 0. Including New Year's Day, the calendar pulses did not share one neighbourhood profile ($P < 0.001$); without it, the three pulses shown still differed ($P = 0.0015$, a test added after fitting). Estimates attach to incident location and are ecological (Supplementary Table S50).

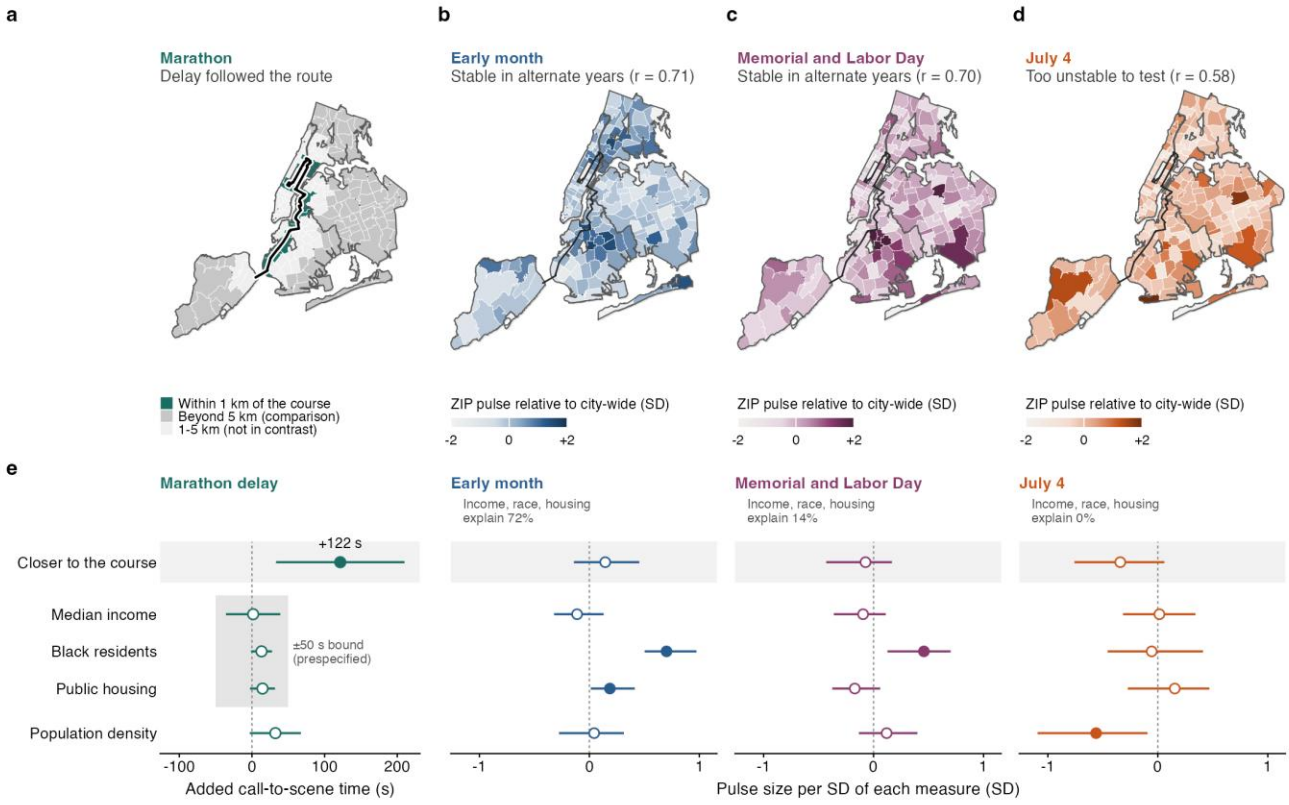

a
Marathon
Delay followed the route
Within 1 km of the course
Beyond 5 km (comparison)
1-5 km (not in contrast)
b
Early month
Stable in alternate years (r = 0.71)
ZIP pulse relative to city-wide (SD)
-2
0
+2
c
Memorial and Labor Day
Stable in alternate years (r = 0.70)
ZIP pulse relative to city-wide (SD)
-2
0
+2
d
July 4
Too unstable to test (r = 0.58)
ZIP pulse relative to city-wide (SD)
-2
0
+2
e
Marathon delay
+122 s
Closer to the course
Median income
Black residents
Public housing
Population density
±50 s bound
(prespecified)
-100
0
100
200
Added call-to-scene time (s)
Early month
Income, race, housing
explain 72%
Memorial and Labor Day
Income, race, housing
explain 14%
July 4
Income, race, housing
explain 0%
-1
0
1
Pulse size per SD of each measure (SD)

Figure 5. The calendar of drug or poisoning calls and how overlapping schedules change holiday estimates. (a) Observed drug or poisoning calls on each day of the year relative to the mean of days 10-24 in the same month and year, averaged over 2012-2023, with daily counts scaled to 24 physical hours. The outline marks days 1-3, and boxes mark fixed-date holidays (New Year's Day, July 4, Halloween, Christmas and New Year's Eve); New Year's Day (+133%) exceeds the colour scale. Memorial Day, Labor Day and Thanksgiving fall on different dates each year and are spread across cells. The panel is descriptive and not adjusted for weekday or weather. (b) Overlap test using the matched designs: the full-day July 4 drug or poisoning excess estimated alone against matched days 7 and 14 days away, the day-4 tail of the month cycle in months other than July, and the July 4 excess beyond that tail, with 95% confidence intervals (Supplementary Table S31, Panel F). The parts combine on the ratio scale, so the excess beyond the tail is the ratio of the first two rows rather than their difference. (c) Holiday day-0 effects against same-weekday days, the same month's ordinary Sunday against same-weekday days, and the holiday beyond the Sunday pattern, from the joint calendar model with year-month-clustered 95% confidence intervals; filled symbols mark intervals that exclude 0. The designs in b and c were fixed before fitting (Supplementary Table S31).

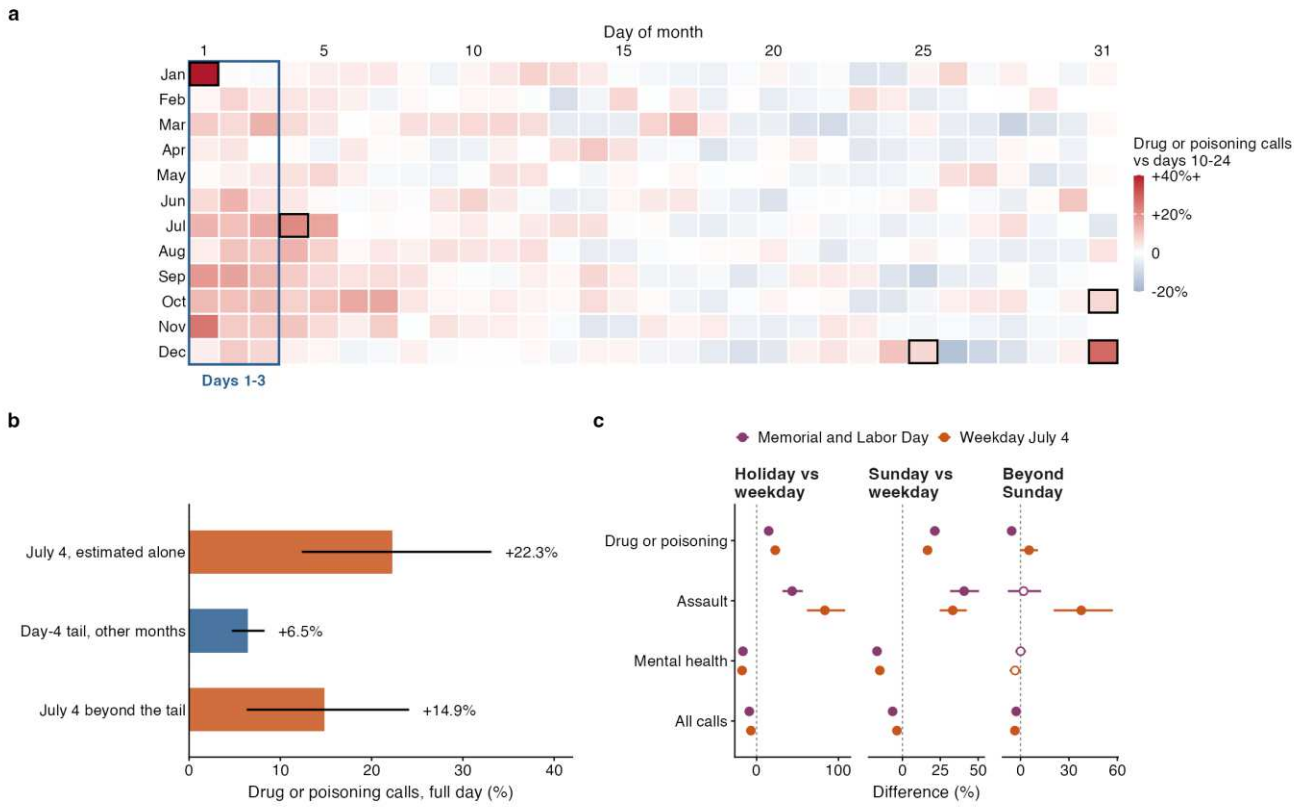


Figure 6. Calls and response times combined as pre-arrival waiting. (a) Change in the number of calls and in mean valid call-to-scene time for each schedule, from its reported matched comparison: early-month days 1-3 against same-weekday days 10-24 (citywide); the July 4 fireworks window (18:00-05:59) against the same hours 7 and 14 days away (citywide); Memorial and Labor Day against the same month's Sunday pattern (citywide); and marathon active operations (07:00-20:59) in route-adjacent ZIPs and United Nations General Assembly weekdays in UN-adjacent ZIPs, each against distant ZIPs and matched control days. Corner labels describe where a schedule's changes fall. Point area is proportional to the absolute change in total waiting, printed beside each point; filled symbols mark schedules whose call and wait intervals both exclude 0. The marathon call change covers the active window (+24%) rather than the whole Sunday (+15.2% in the text), and all wait changes use a common 3,600-second validity ceiling, under which the plotted UN wait change is +68 seconds; with the UN design's 7,200-second ceiling it is +106 seconds, close to the +113 seconds in Fig. 1c. (b) Change in total pre-arrival waiting (incidents multiplied by mean call-to-scene time) relative to the comparison level, divided into parts from more calls at the usual wait, from slower responses for the usual calls, and their interaction; totals and 95% intervals are printed beside each bar. The division is an accounting identity relative to the comparison level, not an attribution of cause. It weights every incident's wait equally regardless of acuity, invalid intervals take the cell mean, and windows, populations and comparisons differ across schedules, so the panels describe a typology rather than a standardized comparison. Totals in incident-hours are in Supplementary Table S46. Placebo rows, shown without intervals, are the 2012 and 2020 fall-back Sundays without a race and the virtual 2020 General Assembly. Intervals are percentile intervals from 2,000 cluster bootstrap draws (month-years, occurrences, races or years); with 10 races and 11 years, the marathon and UN intervals are approximate and narrower than t intervals for the same quantities (marathon wait, 108-202 against 93-210 seconds). The design was fixed after the separate results were known.

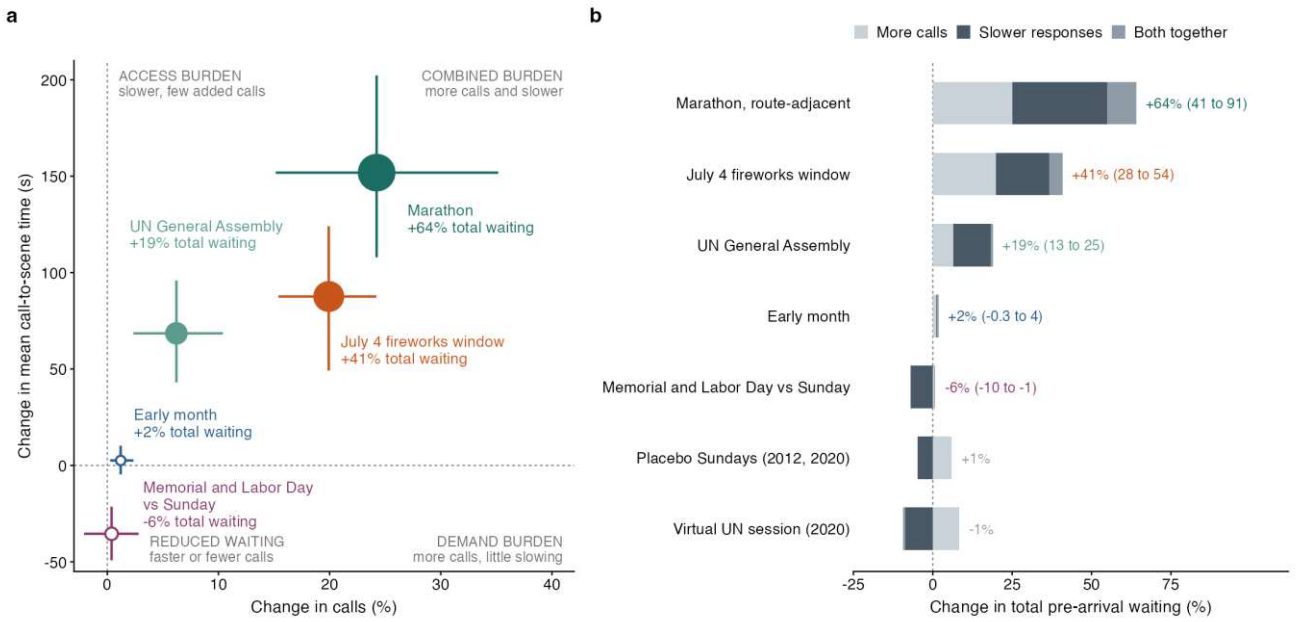
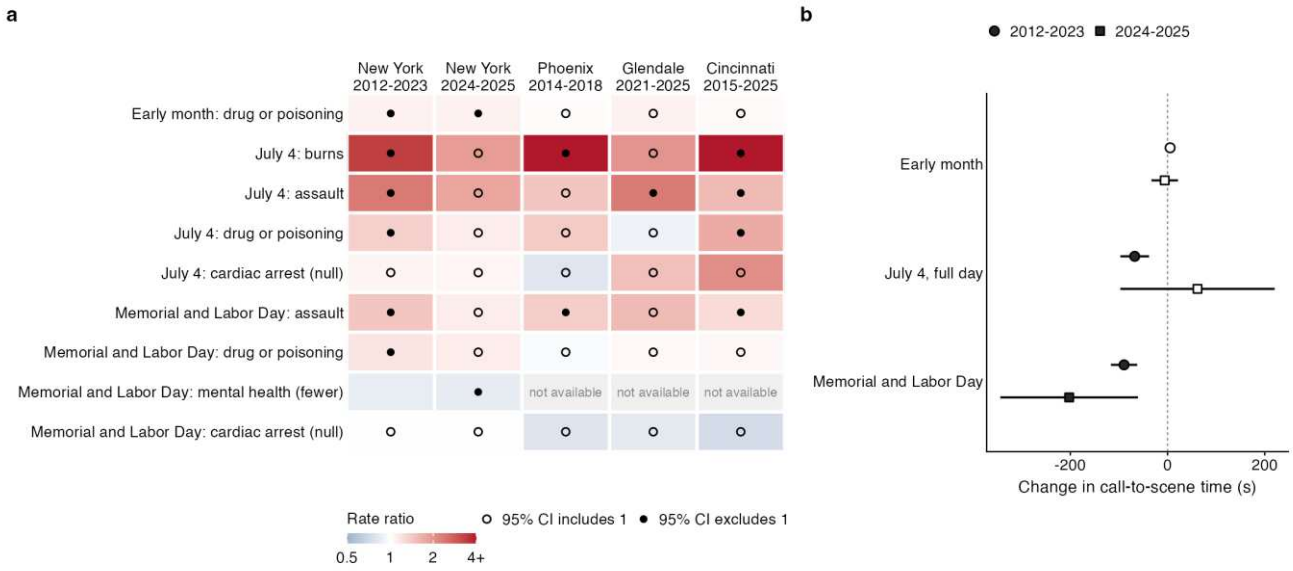


Figure 7. Which patterns recurred. (a) Rate ratios for the seven prespecified schedule-by-endpoint cells and the cardiac-arrest comparison cells, in New York City in 2012-2023 and 2024-2025 and in three external cities, from identical unadjusted matched count contrasts (Supplementary Tables S18 and S24). Colour shows the ratio on a log scale; filled dots mark 95% intervals that exclude 1 and open dots intervals that include 1. As for the external cities, New York 2024-2025 intervals use t references with one fewer degree of freedom than occurrences, because the validation period holds only two July 4 and four Memorial and Labor Day occurrences. No common mental-health category was available outside New York, and the locked ledger stores no interval for the unadjusted 2012-2023 mental-health contrast (the weather-adjusted estimate and its interval are in Extended Data Table 5). Directions, not magnitudes, are compared: city periods differ, and Phoenix and Glendale share a regional code vocabulary. The validation was only partly prospective because 2024 endpoint counts had appeared in an exploratory update; 2025 counts were untouched. The scatter of 2012-2023 against 2024-2025 ratios for all 15 cells, with its Spearman correlation and calibration slope, is in Supplementary Fig. S5d. (b) Citywide call-to-scene contrasts for the three calendar schedules (full holiday day for July 4) in 2012-2023 and 2024-2025, from total-call-weighted city-hour means, with 95% confidence intervals using t references with one fewer degree of freedom than event clusters; filled symbols mark intervals that exclude 0. Intervals and the prespecified decision are reported in Supplementary Table S18.



Tables

Table 1. Independent passenger-trip corroboration during permitted street events. Official high-volume for-hire vehicle records were compared within fixed origin-destination and pickup-hour cells, using baseline control-trip weights and an exposed-minus-unexposed event/control contrast. The primary estimate averages events equally; intervals use 1,999 date-cluster bootstrap draws. The EMS correlation pairs independently estimated event-level contrasts; it does not measure causal mediation. Retained cells cover 23.0% of eligible exposed trip appearances and 37.7% of comparisons. Placebos cover only 50 of the 263 events and use a stricter one-day control-support requirement; the primary estimate within those same 50 events is +70.3 s (95% interval 40.6-100.7). The associations describe selected completed trips, whose routes were not observed. Supplementary Table S26 reports original support and secondary contrasts; Supplementary Fig. S3a and Supplementary Table S28 show the sensitivity using comparison-date weights and a matched placebo.

Comparison	Unit	Estimate	95% interval	Events	Dates
All events	seconds	+53.9	[43.3, 64.7]	263	166
2024	seconds	+49.6	[34.6, 66.4]	109	70
2025	seconds	+57.0	[41.6, 72.4]	154	96
Shifted-date placebo	seconds	-11.0	[-25.9, 2.5]	50	42
EMS-transport association	Pearson r	0.289	[0.155, 0.403]	263	166

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- 02SupplementaryInformation.pdf