ECONOMISTS FOR UKRAINE

WORKING PAPER SERIES

WORKING PAPER No. 38 JUNE 2026

SLEEPLESS NIGHTS, SAFER ROADS? NIGHT-TIME AIR ALERTS AND BEHAVIOURAL RISK AVOIDANCE

Dariia Mykhailyshyna

I thank Monika Pompeo, David Zuchowski, and Xinhui Sun for valuable comments, as well as the audiences at Uppsala University, the Stockholm Institute of Transition Economics (SITE), and the Kyiv School of Economics. I am also grateful to Anna Kovtun for excellent research assistance. All errors are my own.

Economists for Ukraine (Econ4UA)

Website: https://econ4ua.org/ Email: info@econ4ua.org

ABSTRACT

Sleepless Nights, Safer Roads? Night-Time Air Alerts and Behavioural Risk Avoidance

Do salient threats make people take more or less risk? I study this question using wartime air alerts in Ukraine, a recurring shock that is threatening, disruptive, and often occurs at night. Combining time-stamped alert records with daily administrative traffic-accident data from Lviv, I find that each additional hour of overnight alert exposure is associated with about 10 percent fewer next-day accidents. The decline is concentrated in afternoon and evening travel, weekends, central areas, and accident categories linked to risky driving, while fatigue-related accidents do not increase. These patterns are consistent with behavioural risk avoidance: after a salient threat, people appear to reduce or retime discretionary travel and drive more cautiously. Evidence from thirteen other Ukrainian cities suggests that this response is strongest where alerts are relatively rare and weaker where alerts are routine. Salient threats can therefore reduce risky behaviour, but the response appears to depend on continued salience.

JEL CLASSIFICATION: R41, D91, D74

KEYWORDS: Traffic accidents, risk preferences, risk avoidance, salient threat, behavioural adaptation, air alerts, conflict

Dariia Mykhailyshyna
Kyiv School of Economics
dmykhailyshyna@kse.org.ua

Sleepless Nights, Safer Roads?
Night-Time Air Alerts and Behavioural Risk Avoidance

Dariia Mykhailyshyna*

June 19, 2026

Abstract

Do salient threats make people take more or less risk? I study this question using wartime air alerts in Ukraine, a recurring shock that is threatening, disruptive, and often occurs at night. Combining time-stamped alert records with daily administrative traffic-accident data from Lviv, I find that each additional hour of overnight alert exposure is associated with about 10 percent fewer next-day accidents. The decline is concentrated in afternoon and evening travel, weekends, central areas, and accident categories linked to risky driving, while fatigue-related accidents do not increase. These patterns are consistent with behavioural risk avoidance: after a salient threat, people appear to reduce or retime discretionary travel and drive more cautiously. Evidence from thirteen other Ukrainian cities suggests that this response is strongest where alerts are relatively rare and weaker where alerts are routine. Salient threats can therefore reduce risky behaviour, but the response appears to depend on continued salience.

JEL Codes: R41, D91, D74

*Kyiv School of Economics. Email: dmykhailyshyna@kse.org.ua. I thank Monika Pompeo, David Zuchowski, and Xinhui Sun for valuable comments, as well as the audiences at Uppsala University, the Stockholm Institute of Transition Economics (SITE), and the Kyiv School of Economics. I am also grateful to Anna Kovtun for excellent research assistance. All errors are my own.

Keywords: Traffic accidents, risk preferences, risk avoidance, salient threat, behavioural adaptation, air alerts, conflict

1 Introduction

How does exposure to a salient threat change risky behaviour? A large body of work finds that living through violence or disaster makes people more cautious and risk-averse (Cahlíková and Cingl, 2017; Cameron and Shah, 2015; Dohmen et al., 2016); while another strand of literature finds that such exposure raises risk tolerance (Voors et al., 2012; Hanaoka et al., 2018). These contradictory findings can also be found with regard to the behaviour on the road: fatal accidents have been found to rise after terror attacks (Stecklov and Goldstein, 2004), yet salient warnings of danger are widely thought to induce caution (Anand, 2025; Ferris and Newburn, 2017). Whether a salient threat leaves people driving more or less dangerously is therefore an open empirical question, which is not easily answered because threats rarely recur at a high enough frequency to allow an empirical examination.

This paper studies a setting where, unfortunately, such a threat does recur with the needed frequency: wartime air-raid alerts in Ukraine. When aerial threats are detected, sirens sound across a region, often at night, signalling a potential threat to residents' lives on an irregular and unpredictable schedule. The alerts also cost sleep, a channel that, as the sleep-and-driving literature shows, raises crash risk (Smith, 2016; Bünnings and Schiele, 2021; Flynn et al., 2025). I treat sleep disruption as a competing explanation and show some suggestive evidence that it does not drive the results. Because the threat is external in origin and the timing of any given night's alert is unpredictable locally, the alerts generate repeated, plausibly exogenous variation. I ask whether this variation is followed by more or fewer accidents the next day, and examine the mechanisms driving the results.

I begin with Lviv, a major city in Western Ukraine, with the richest administrative microdata, combining daily counts of police-recorded accidents over 2022–2024 with time-stamped alert logs. I estimate Poisson pseudo-maximum-likelihood (PPML) models of next-day accident counts on overnight alert exposure, absorbing day-of-week, month, and year fixed effects. Identification rests on the unpredictable timing of alerts: conditional on the calendar, whether a given night is disrupted is as good as random with respect to the next day's driving. Three diagnostics support this. First, alert timing is poorly predicted by observable local conditions (pseudo-R^2 of 0.10 with the full set of local controls). Second, lead tests show no anticipation. Third, observable covariates are balanced across alert and non-alert nights.

Overnight alert exposure is associated with fewer next-day accidents. Each additional hour of alert during the night window (22:00–05:59) corresponds to roughly 10 percent fewer full-day accidents (06:00–21:59), with the reduction concentrated in the afternoon and evening (about 15 percent in each window) and no detectable change in the morning. The morning null is consistent with commuting trips that are hard to re-time, whereas the afternoon and evening capture more discretionary travel. This null is specific to Lviv's injury-only records: the all-accident cross-city data in Section 6 reveal a morning decline concentrated in the minor, property-damage collisions that Lviv records do not include. The estimates are stable across a wide set of local controls.

Next, I explore potential mechanisms. The results point to behavioural risk avoidance operating on two margins. On the *extensive margin*, the patterns are consistent with fewer or retimed discretionary trips. This is supported by three suggestive tests. The reduction is far larger on weekends (38 percent) than on weekdays (5 percent); it concentrates in the city centre (19 percent) rather than the periphery (3.5 percent); and it dissipates within a day, with a suggestive positive rebound, as if trips were postponed rather than forgone.

On the *intensive margin*, those who do drive appear to drive more carefully. Police-recorded accident causes show that a decrease in risky behaviours drives the results: accidents attributed to risky behaviour, such as speeding and improper overtaking, fall by about 21 percent, while accidents attributed to fatigue do not rise. A second test uses parking fines as an imperfect proxy for the level of traffic. Fine volumes are essentially unchanged after alert nights even as accidents fall, so the accident-per-fine ratio declines by about 25 percent, and enforcement-timing checks indicate that this is not an artifact of how tickets are issued.

The protective response is not constant over time. It is strongest in 2022, the first year of the full-scale invasion, with a full-day reduction of about 27 percent, and fades to statistical insignificance in 2023 and 2024. One potential interpretation is adaptation: residents respond most when alerts are novel and frightening, and less as alerts become routine. But the same pattern could instead reflect something particular to the first year of the war. Distinguishing the two requires variation in how routine alerts are while holding the first-year-of-war confound fixed.

That variation exists across cities. I extend the analysis to thirteen other Ukrainian cities, for which accident data could be obtained. These cities lack the auxiliary details that the Lviv mechanism tests rely on, so while I can use the data from them to show the robustness of Lviv results, as well as to test the adaptation hypothesis, I am not able to fully replicate the mechanism decomposition. Alert frequency varies sharply across Ukraine, from cities under alert on less than 15 percent of nights to cities under alert almost nightly. This discrepancy is driven by wartime geography rather than local road conditions. The protective effect tracks that frequency: it is present where alerts are infrequent and absent where they are frequent. Pooled across these cities, a disrupted night is associated with about 4 percent fewer next-day accidents, and the response is concentrated where alerts are rare. This provides additional suggestive evidence in support of adaptation: cities that experience more frequent air alerts have adapted faster and thus stopped adjusting their behaviour in response to an alert.

Adaptation patterns can also be observed when examining stretches of repeated air alert nights. While the response persists from the first night of a run through the second and third-or-later nights with only mild attenuation, the cumulative alert load carries no independent effect once the immediate night is accounted for. A mechanical-fatigue account predicts the opposite, that accumulated sleep debt should matter most. Instead, the response is keyed to the immediate, salient night and decays as alerts lose novelty. Taken together, these three timescales—years, cities, and consecutive nights—point the same way: a behavioural response that erodes as alerts become an ordinary part of life.

This paper contributes to three strands of literature. First, I contribute to work on how exposure to threat moves risk-taking, where the evidence is mixed: studies find greater risk aversion after some natural disasters and crises (Cameron and Shah, 2015; Dohmen et al., 2016) but greater risk tolerance after others (Hanaoka et al., 2018), and, for exposure to violence specifically, effects in both directions have been recorded (Voors et al., 2012; Moya, 2018; Jakiela and Ozier, 2019). That evidence relies mainly on surveys and lab or hypothetical-choice measures, often collected well after the shock. I add to this literature with a high-frequency behavioural test, using actual high-stakes driving outcomes the day after a salient threat. Second, I also contribute to work on civilian behaviour under threat. Becker and Rubinstein (2011) argue that fear, rather than the objective probability of harm, drives the behavioural response to terrorism, and that the response is amplified when attacks are most salient. Stecklov and Goldstein (2004, 2010) also look at traffic accidents and find that fatal accidents in Israel spike by about 35 percent in the days after terror attacks, with no habituation over four years. My study shows the opposite: accidents fall rather than rise, and the response fades rather than persists, consistent with the difference between advance warning and the acute trauma of a realized attack. The fading response connects to declining civilian responsiveness to Ukrainian alerts (Van Dijcke et al., 2023) and complements evidence on other wartime behaviours (Martinez et al., 2025; Klymak et al., 2025).

Third, I contribute to the literature on driving behaviour. The sleep-and-road-safety literature treats lost sleep as unambiguously harmful for driving; I document a case in which sleep loss is instead followed by fewer accidents, because it arrives bundled with a threat that induces caution. Unlike warning-based studies in which advisories about hazardous conditions reduce crashes through a targeted protective channel (Anand, 2025; Ferris and Newburn, 2017), the warning here concerns a danger unrelated to driving, so the safety gain reflects a general shift in caution. Relative to estimates of how transitory shocks move road safety, including pollution (Shr et al., 2023), allergens (Danagoulian and Deza, 2025), fasting (Gulek, 2024), and religious observance (Romem and Shurtz, 2016), I add a very different shock, a war-related threat operating through risk perception rather than physiology.

The remainder of the paper proceeds as follows. Section 2 describes the setting and data. Section 3 presents the empirical strategy. Section 4 reports the main results for Lviv. Section 5 examines mechanisms. Section 6 examines how the effect generalizes across cities and how it adapts over time. Section 7 presents robustness checks, and Section 8 concludes.

2 Setting and Data

2.1 Air Alerts in Ukraine

On February 24, 2022, Russia launched a full-scale invasion of Ukraine, with its troops crossing into Ukraine all along the Ukrainian-Russian border and missiles striking targets across the country. Since then, while Ukraine has resisted the Russian invasion on the ground, Russian missile and drone attacks have continued to target Ukrainian cities on a regular basis. In order to warn residents of potential danger, since the very first days of the war, the Ukrainian government has implemented a system of air alerts, which are activated when potential aerial threats (missiles or drones) are detected, warning residents to seek shelter. Alerts are managed at the regional level and announced through official channels, including mobile applications, sirens, and media. Because the underlying threats are external, the frequency and timing of alerts are driven by exogenous factors rather than by conditions in any one city. Eastern regions face more frequent and severe threats, while western cities including Lviv still experience regular alerts, particularly at night when long-range attacks are more common.[1]

As mentioned previously, the frequency of alerts varies throughout the country. In the cities I study, the share of nights with any alert ranges from roughly one in eight in the

[1]While night-time air alerts occur on fewer days (17 percent of days have an air alert at night vs. 37 percent in daytime), they are disproportionately concentrated within a shorter window (8 hours vs. 16 hours), making them more salient per unit of time. Long-range missile and drone attacks also tend to be launched at night to exploit longer flight paths under cover of darkness, and night alerts are particularly disruptive, as residents are awoken by the loud sound of sirens and advised to move to shelter.

far-western oblasts to almost every night in the south and east. This cross-city variation in alert frequency is the basis for the generalization and adaptation analysis in Section 6.

2.2 The Lviv Setting

Lviv, located in western Ukraine near the Polish border, provides an informative setting for several reasons. First, the city experiences substantial alert variation without the confounding direct damage that affects frontline areas. Second, the city has well-functioning administrative data systems that continued operating throughout the war, and its accident records are unusually rich, combining the largest accident counts in the sample with geocoordinates, officer-coded causes, severity, and a parking-fine activity proxy that the mechanism tests require. Other Ukrainian cities also provide accident microdata in response to freedom-of-information (FOI) requests, but many of them have far fewer accidents and do not have these auxiliary variables, which is why Lviv anchors the analysis while the other cities support the cross-city comparison. Third, Lviv is one of the largest cities in Ukraine and is a popular tourist destination. The urban environment features both discretionary destinations (the historic centre, restaurants, shopping) and necessary travel corridors (commuting routes, industrial areas), allowing heterogeneity analysis.

An important caveat is that Lviv's population changed during the study period, as it received a large number of internally displaced persons (IDPs) from Eastern and Central Ukraine, while some residents moved abroad. These demographic shifts may affect both the population at risk and the composition of drivers, though month and year fixed effects absorb gradual changes in population levels.

2.3 Data Sources

I combine several data sources. First, traffic accident records from the Lviv regional division of the National Police of Ukraine contain the universe of reported accidents with injuries or fatalities in Lviv city, including timestamp, location coordinates, accident type, severity, and official cause classification. It is important to note that the accident data covers reported accidents with injuries or fatalities only; property-damage-only incidents are not included in the police database. This means the analysis captures relatively serious accidents but may understate the overall effects.

Second, air alert records from the official Ukrainian alert system, collected by Klymenko (2024) provide start and end times for each alert at the regional level. Third, hourly weather data from the Open-Meteo Historical Weather API covers temperature, precipitation, wind speed, humidity, and cloud cover for Lviv.[2] Fourth, parking violation records from the Lviv Municipal Guard cover all parking fines issued in the city, with timestamp and location, providing a proxy for traffic activity.

Fifth, daily PM2.5 air quality data from the SaveEcoBot citizen-science network provide a control for pollution-related driving impairment.[3] Sixth, I use conflict event data from the VIINA database (Zhukov and Ayers, 2023) recording missile and drone attacks on Lviv city. I supplement these sources with institutional calendars marking public holidays, school sessions, and curfew regulations.

The main analysis sample covers February 26, 2022 (the date air alert data became systematically available) through December 31, 2024, yielding 1,040 days. I aggregate accidents to the day level, constructing counts for four time windows: morning (06:00–11:59), afternoon (12:00–17:59), evening (18:00–21:59), and full day (06:00–21:59). The night window (22:00–05:59) is used primarily to construct the treatment rather than as an outcome, because Ukraine's martial-law curfew prohibits being outside during most of those hours, except in emergencies. At the same time, as a robustness check, I also explore contemporaneous effects of the night-time air alert on night-time traffic accidents.

The main treatment variable of interest is the total night alert exposure in hours, calculated as the total alert duration in the Lviv region during the preceding night. An

[2] Open-Meteo Historical Weather API, `https://open-meteo.com/en/docs/historical-weather-api`.

[3] SaveEcoBot (`https://www.saveecobot.com/en`) aggregates data from citizen-operated air quality sensors across Ukraine.

alert beginning before 22:00 and ending after 22:00 contributes only the portion after 22:00; similarly, alerts extending past 06:00 are truncated.

Table 1 reports summary statistics. The sample averages 1.7 accidents per day during daytime hours. Night alert exposure is on average 11.5 minutes per night, with 17 percent of days experiencing any night alert. Alert duration conditional on occurrence averages approximately 67 minutes, with substantial variation (standard deviation of 32 minutes unconditionally).

Table 1: Summary Statistics

Variable	N	Mean	SD	Min	Max
Panel A: Treatment					
Night Alert Duration (minutes)	1,040	11.46	32.28	0.00	258.72
Any Night Alert (0/1)	1,040	0.17	0.38	0.00	1.00
Panel B: Outcomes (Daily Accident Counts)					
Morning Accidents (06–11)	1,040	0.43	0.68	0	4
Afternoon Accidents (12–17)	1,040	0.75	0.87	0	5
Evening Accidents (18–21)	1,040	0.52	0.73	0	3
Full Day Accidents (06–21)	1,040	1.70	1.38	0	7
Panel C: Controls					
Temperature (C)	1,040	11.71	8.67	−13.8	30.2
Precipitation (mm)	1,040	1.59	3.31	0.0	39.3
Wind Speed (km/h)	1,040	12.96	5.37	2.2	34.4
Humidity (%)	1,040	72.05	13.24	34.1	98.9
Cloud Cover (%)	1,040	68.53	30.09	0.0	100.0
Weekend (0/1)	1,040	0.29	0.45	0.0	1.0
Public Holiday (0/1)	1,040	0.03	0.17	0.0	1.0
Daylight Hours	1,040	12.42	2.80	8.1	16.3
Any Lviv City Attack (0/1)	1,040	0.14	0.35	0.0	1.0
Daytime Alert Duration (minutes)	1,040	27.61	52.05	0.0	413.0
School Session (0/1)	1,040	0.72	0.45	0.0	1.0
Curfew Duration (hours)	1,040	5.32	0.71	5.0	8.0

Notes: Sample period: February 26, 2022 to December 31, 2024 (1,040 days). Night alert duration measures total minutes of air alert exposure during the night window (22:00–05:59) preceding the outcome day. Accidents are counts of traffic accidents with injuries in Lviv city.

Figure 1 displays the spatial distribution of accidents in Lviv over the sample period. Accidents concentrate in the city centre and along major roads, with notable clustering around the historic district, main intersections, and commercial areas. This spatial pattern is used for heterogeneity analysis in the later sections.

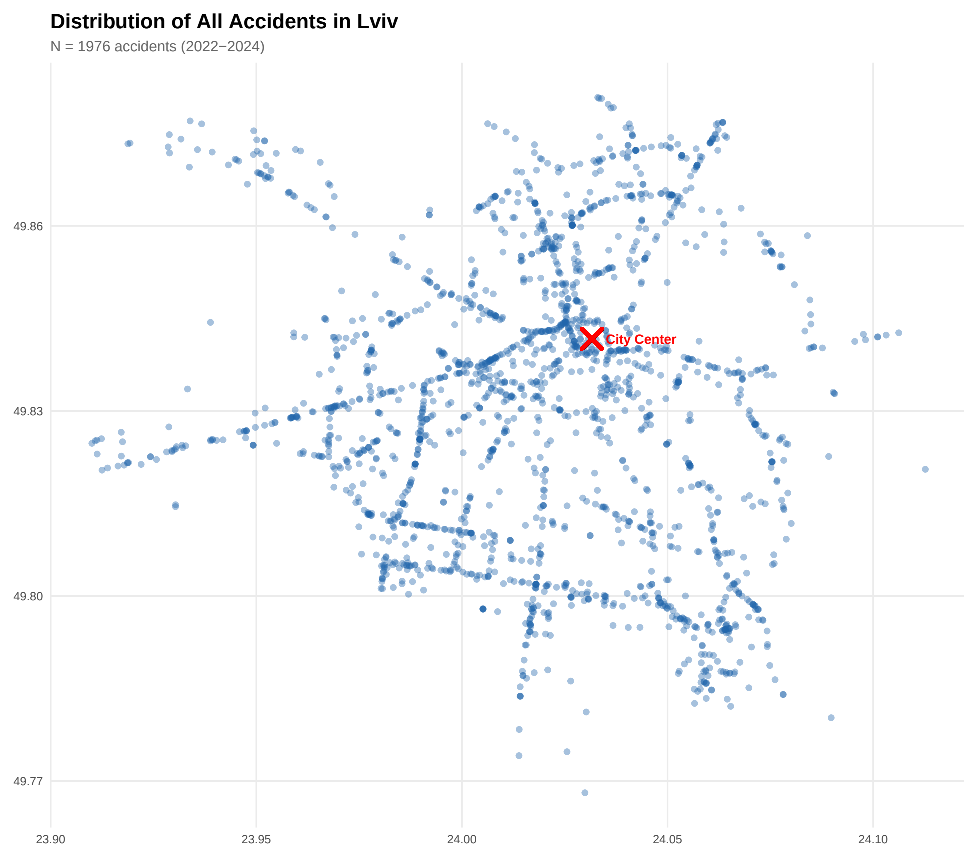

Figure 1: Spatial Distribution of Traffic Accidents in Lviv, 2022–2024
Notes: Each point represents a traffic accident with injuries or fatalities reported to the National Police. The sample covers February 2022 through December 2024 ($N = 1,976$ accidents with valid geocoordinates). Accidents concentrate in the city centre (historic district, commercial areas) and along major roads.

2.4 Cross-City Data

To assess how the Lviv findings generalize, I assemble comparable data for thirteen additional Ukrainian cities: Chernihiv, Chernivtsi, Dnipro, Lutsk, Mykolaiv, Odesa, Poltava, Rivne, Sumy, Ternopil, Uzhhorod, Vinnytsia, and Zaporizhzhia. As for Lviv, the records are accident-level extracts from the National Police accident register, obtained through FOI requests to the regional police divisions, with officer-recorded timestamps, accident types, injury outcomes, and cause classifications.[4] The set comprises the cities for which the requests yielded register extracts with the time-of-day information the window analysis requires; it includes neither of Ukraine's two largest cities, because comparable microdata for Kyiv could not be obtained and the available Kharkiv records lack time stamps. Data availability is an administrative matter of which regional divisions fulfilled the requests, and the resulting dataset spans the full range of wartime exposure, from rarely alerted western cities to heavily targeted eastern and southern ones; still, I cannot rule out that availability correlates with unobserved features of local record-keeping.

For each city I construct daily accident counts by time window and night alert exposure exactly as for Lviv, using the same alert source. The cross-city panels cover February 2022 through December 2025, one year longer than the Lviv sample, whose accident records end in December 2024. Two features distinguish these datasets from Lviv's. First, some of them are thinner: the other western cities in particular have substantially fewer accidents and do not include parking-fine records that the Lviv mechanism tests rely on, so they support the reduced-form cross-city comparison but not the margin decomposition. Second, unlike the Lviv injury data, they include property-damage accidents. I therefore harmonize the cross-city outcome to all reported accidents for the main analysis and, as a robustness check, to the injury-and-fatal subset comparable to Lviv. Alert frequency varies widely across these cities, which is the source of identifying variation for the cross-city analysis in Section 6.

For each city I also record, from the VIINA conflict-event database (Zhukov and Ayers, 2023), how often an alert coincides with a nearby strike: $P(\text{strike} \mid \text{alert})$ is the share of alert days on which a strike occurred within 25 kilometres of the city centre—airstrikes and drone strikes for every city, and additionally artillery for the six frontline cities within shelling range (Dnipro, Mykolaiv, Odesa, Sumy, Zaporizhzhia, and Chernihiv). It is near zero in the western cities and rises to between 12 and 41 percent in the most heavily attacked eastern and southern cities, providing a descriptive measure of how directly each city is exposed to

[4] Not all of this information is available for all cities. For instance, Chernihiv and Mykolaiv do not have exact geolocation of an accident.

attack (Table 6).

3 Empirical Strategy

3.1 Specification

I estimate Poisson pseudo-maximum likelihood (PPML) models of the form:

$$E[Y_{d,w}|\mathbf{X}_d] = \exp\left(\beta \cdot \text{NightAlertHours}_d + \boldsymbol{\gamma}'\mathbf{FE}_d + \boldsymbol{\delta}'\mathbf{X}_d\right) \tag{1}$$

where $Y_{d,w}$ is the count of accidents on day d in time window w, NightAlertHours_d is total duration in hours of air alerts in Lviv region during the preceding night, $\mathbf{FE}_d$ includes fixed effects for day-of-week, month, and year with coefficient vector $\boldsymbol{\gamma}$, and $\mathbf{X}_d$ includes time-varying controls, used in robustness specifications with coefficient vector $\boldsymbol{\delta}$.

The coefficient β is interpreted as the semi-elasticity: the proportional change in expected accidents from a one-hour increase in night alert exposure. I report percentage effects calculated as $\exp(\beta) - 1$, the percentage change in expected accidents per hour of exposure; the corresponding incidence rate ratio is $\exp(\beta)$. Standard errors are heteroskedasticity-robust.[5]

3.2 Identification

My main identifying assumption is that night alert timing is not correlated with unobserved determinants of next-day accidents, conditional on controls. In other words, conditional on day-of-week, month, and year, the timing and duration of night alerts in Lviv are as good as random with respect to local factors affecting traffic safety.

I present several arguments in favour of this assumption. First, air alerts are triggered

[5] I tested for nonlinearity by adding a quadratic term (NightAlertHours^2) to the baseline specification; the squared coefficient is small and insignificant across all time windows, confirming that the linear specification is adequate.

when aerial threats are detected en route, or when Russian fighter planes are detected on the launching positions. The precise timing of any given night's alert is unanticipated by Lviv residents – they cannot predict whether tonight there will be an attack on Ukraine, nor whether it will target Lviv specifically, nor when it will start or how long it will last. While some seasonality exists, in my analysis, this is absorbed by month and year fixed effects. Secondly, as argued in more detail in Klymak et al. (2025), the effectiveness of Russian missile attacks relies on their unpredictability: if it were possible to anticipate in advance where and when Russia would strike, it would significantly reduce the damage that such attacks intend to cause. Thirdly, additional diagnostic tests, presented in Section 7, support this assumption. Section 4.2 provides a visual diagnostic: in an event-study specification, pre-treatment leads of alert exposure are jointly insignificant, so alert nights do not follow unusual accident days.

It is useful to walk through the three canonical threats to identification in turn. *Unobserved heterogeneity*: a confounder would have to vary at the daily level within calendar cells and move both alert timing and next-day driving conditions. The fixed effects absorb the calendar; observable candidates – weather, holidays, daylight, local strikes, curfew changes – are balanced across alert and non-alert nights (Section 7) and leave the estimate unchanged when included (Table 9).

Reverse causality: traffic conditions in Lviv cannot plausibly influence the timing of Russian missile launches, and the data provides support for this – tomorrow's alerts do not predict today's accidents (Table A6), and alert timing is poorly predicted by local observables (Table A5).

Measurement error: the treatment is the duration of oblast-level alerts, a proxy for the disruption residents actually experience. Some residents sleep through short alerts or mute notification apps, while others lose sleep beyond the alert window, so true exposure is measured with error. If this is the case, it attenuates the estimated coefficient towards zero, making the estimates conservative lower bounds on the behavioural response to experienced disruption.

4 Results

4.1 Main results

Table 2 reports the main estimates. Each column corresponds to a different time window: morning (06:00–11:59), afternoon (12:00–17:59), evening (18:00–21:59), and full day (06:00–21:59). The mechanism predicts no response in the largely non-discretionary morning commute and a response in afternoon and evening travel, so I read the pattern across columns as one test of that prediction rather than as four separate hypotheses.

Column (1) shows that morning traffic accidents are not significantly affected by night alert exposure. This null result is consistent with morning travel being largely non-discretionary (commuting to work, school runs), limiting the scope for trip suppression; it does not rule out a response that appears only in accident types these injury-only records cannot capture. Indeed, the cross-city data of Section 6, which include property-damage collisions, show a morning decline concentrated precisely in such minor accidents, while morning injury accidents remain flat there as well.

Columns (2) and (3) present evidence for reductions in afternoon and evening accidents after the air alert during the preceding night. Afternoon accidents decline by 14.9 percent per hour of night alert, while evening accidents drop by 15.1 percent. These results are consistent with reductions in discretionary travel: in the afternoons and evenings, people tend to shop, dine, and engage in other social activities where behavioural adjustment is feasible.

Finally, column (4) shows the main results of the effect of night-time air alerts on the total number of accidents the next day. Accidents decline by 9.9 percent per hour of night alert exposure. Relative to a baseline mean of 1.7 accidents per day, this implies roughly 0.17 fewer accidents per hour of alert, or approximately one avoided accident per six hours of cumulative night exposure. Because the per-hour coefficient is a semi-elasticity, it is also useful to state the effect of typical exposures: the average alert night involves about 67 minutes of alerts, so a typical alert night is followed by roughly 11 percent fewer accidents, or 0.19 fewer accidents. A one-standard-deviation increase in exposure (about half an hour) implies roughly 5 percent fewer accidents. Because this dataset captures only traffic accidents with injuries or lethal outcomes, these results may represent a lower bound on the overall effect of night-time air alerts on accidents.

Panel B of Table 2 repeats the specification with a binary indicator for any night alert. The full-day point estimate is nearly identical (about a 9 percent reduction) but no longer significant, and the afternoon effect similarly weakens, while the evening reduction remains significant. This pattern indicates that the response scales with how long the night is disrupted; I therefore treat the continuous specification as primary.

To gauge this magnitude, it helps to compare it with related shocks that impair drivers. The spring transition to daylight saving time, which costs about an hour of sleep, raises fatal crashes by roughly 6 percent (Smith, 2016), with some estimates as high as 8 to 10 percent (Bünnings and Schiele, 2021); Ramadan fasting raises the probability of causing an accident by about 25 percent (Gulek, 2024). In each case a physiological burden makes driving more dangerous. Here, by contrast, a comparable night-time disruption is followed by about 10 percent fewer accidents, so the threat channel does not merely offset the expected fatigue effect but reverses its sign. The contrast with the conflict-behaviour literature is sharper still: Stecklov and Goldstein (2004, 2010) find that fatal accidents rise by about 35 percent in the days after a terror attack, whereas here both light and severe injury accidents decline; the fatal series, at roughly two accidents per month, is too thin to be informative (Table A8).

Table 2: Effect of Night-Time Alert Exposure on Next-Day Traffic Accidents

	(1) Morning (06:00–11:59)	(2) Afternoon (12:00–17:59)	(3) Evening (18:00–21:59)	(4) Full Day (06:00–21:59)
Panel A: Continuous treatment (night alert hours)				
Night Alert Hours	0.034	−0.162**	−0.163*	−0.104**
	(0.085)	(0.073)	(0.087)	(0.045)
IRR (% change)	[+3.4%]	[−14.9%]	[−15.1%]	[−9.9%]
95% CI	[0.88, 1.22]	[0.74, 0.98]	[0.72, 1.01]	[0.82, 0.98]
Panel B: Binary treatment (any night alert)				
Any Night Alert	0.052	−0.094	−0.247*	−0.097
	(0.132)	(0.096)	(0.129)	(0.065)
IRR (% change)	[+5.3%]	[−9.0%]	[−21.9%]	[−9.3%]
95% CI	[0.81, 1.36]	[0.75, 1.10]	[0.61, 1.01]	[0.80, 1.03]
Mean accidents/day	0.43	0.75	0.52	1.70
Observations	1,040	1,040	1,040	1,040
Day-of-week FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes

Notes: Poisson pseudo-maximum likelihood (PPML) estimates. Sample: Lviv, February 2022–December 2024 (1,040 days). Dependent variable is the daily count of traffic accidents with injuries in each time window. Panel A uses the continuous treatment, total hours of air alert exposure during 22:00–05:59 preceding the outcome day; Panel B uses a binary indicator for any night alert. IRR = Incidence Rate Ratio, calculated as $\exp(\beta) - 1$: in Panel A the percentage change per additional alert hour, in Panel B the change relative to a night with no alert. The continuous specification is the primary one, since the response scales with the duration of disruption. Heteroskedasticity-robust standard errors in parentheses. * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

4.2 Dynamic Effects and Pre-Trends

Figure 2 presents a local projections event study following Jordà and Taylor (2025). Each point is the coefficient from a separate OLS regression in long differences ($y_{t+h} - y_{t-1}$ on NightAlertHours_t) at horizon h, with lag augmentation as recommended by Montiel Olea and Plagborg-Møller (2021); for $h > 0$, regressions control for the intermediate treatment path.

The pre-treatment side of the figure is the visual counterpart of the identifying assumption in Section 3.2: if alert nights followed unusual accident days, the lead coefficients would differ from zero. No lead is individually significant at the five percent level, and they are jointly insignificant (joint $p = 0.84$), indicating that alert timing does not respond to local accident dynamics. The post-treatment side shows that the response is concentrated at the contemporaneous horizon: the effect at $h = 0$ is a -8 percent reduction, in line with the baseline estimates in Table 2, and returns towards zero thereafter. The point estimates at longer horizons oscillate in sign but remain noisy and do not trace a systematic dynamic path. I therefore treat this event study as an identification diagnostic that provides evidence for the absence of pre-trends, rather than as direct evidence on the behavioural mechanism.

A distributed lag model, regressing today's accidents on alerts during the preceding night and the three nights before that, shows a consistent pattern (Appendix Table A1): the effect concentrates at lag zero, fades by lag one, and turns positive and significant at lag two (9.6 percent). The positive lag-two coefficient is consistent with a portion of suppressed trips being rescheduled a couple of days later. Overall, the effect does not persist or grow, which is consistent with an immediate behavioural response rather than with cumulative fatigue, which would predict persistent or increasing effects. Section 6 returns to the question of air alerts on consecutive nights with the better-powered cross-city sample, where the effect appears already on the first night of a run and does not build up over a sequence of disrupted nights.

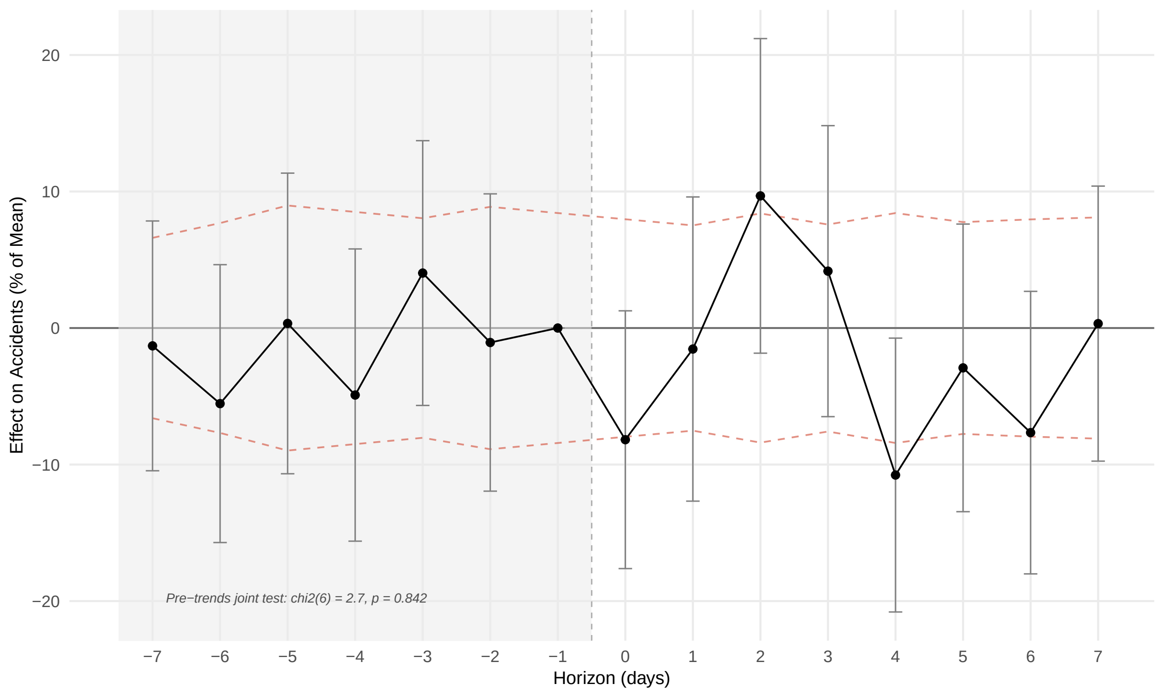

Figure 2: Event Study: Dynamic Effect of Night Alert Hours on Full-Day Accidents

Notes: Local projection estimates following Jordà and Taylor (2025). Each point reports the OLS long-difference coefficient at horizon h: $y_{t+h} - y_{t-1}$ regressed on night alert hours, with lag augmentation (Montiel Olea and Plagborg-Møller, 2021) and day-of-week, month, and year fixed effects. For $h > 0$, regressions control for the intermediate treatment path. Error bars show pointwise 95 percent confidence intervals (heteroskedasticity-robust). Dashed red lines are significance bands (Jordà and Taylor, 2025). The shaded region marks pre-treatment horizons; jointly insignificant leads support exogeneity (joint $p = 0.84$). $h = -1$ is the reference period. Effects expressed as percent of the outcome mean.

5 Mechanisms

The traffic accident reduction could operate on two margins: an extensive margin (reduced travel) and an intensive margin (safer driving conditional on travel). I present some suggestive evidence on both, beginning with the extensive margin (discretionary trip reduction), then turning to the intensive margin through cause composition analysis and a parking-fine proxy for conditional risk.

5.1 Extensive Margin: Discretionary Trip Reduction

If people reduce travel after the nights with air alerts, the effects should be larger when and where travel is more discretionary, and thus can be easily cancelled or postponed. Table 3 tests this prediction.

In Panel A, I split the sample and separately examine the effect of night-time air alerts on next-day traffic accidents on weekends and weekdays. Weekend travel is predominantly discretionary, while weekday travel includes more necessary commuting for work and school runs. While on the weekends the night-time air alerts reduce the traffic accidents by 38.4 percent, on the weekdays the effect of air alerts is not significant (although still negative). This provides the first piece of suggestive support in favour of discretionary travel reduction.

Panel B compares accidents near versus far from the city centre. Central Lviv contains the historical district, main shopping areas, and restaurants, destinations that are largely discretionary. Peripheral areas include residential neighbourhoods and industrial zones where travel is more often necessary. The full-day effect is –18.7 percent for accidents within 3 km of the centre, while not being significant for accidents further than 3 km from the city centre. A finer decomposition by area type (Appendix Table A2) confirms this pattern: the city centre shows the largest reduction, while residential areas show no significant effect. Both the weekend and central-area contrasts also hold under a binary any-alert specification.

Figure 3 visualizes this spatial heterogeneity using a grid-based analysis. The city is divided into approximately 1.5 km × 1.3 km cells, and separate regressions are estimated for each cell with at least 30 accidents. Risk avoidance effects (blue) concentrate in the central and southern areas containing commercial and entertainment destinations, while null or positive effects (red) appear in peripheral residential areas. The positive peripheral estimates could reflect traffic redistribution, with some discretionary trips shifting from the centre to the periphery (thus locals choose activities closer to home, rather than driving all the way to the city centre) or simply estimation noise in grid cells with lower accident counts. Given the imprecision of cell-level estimates, these red areas should be interpreted with caution. Overall, however, the effect of air alerts on traffic accidents is predominantly negative, and the observed pattern is consistent with the extensive margin operating primarily through reductions in discretionary travel to central destinations.

Table 3: Extensive Margin: Heterogeneity by Trip Discretion

	(1) Weekend	(2) Weekday
Panel A: Weekend vs. Weekday (Full Day)		
Night Alert Hours	−0.485***	−0.051
	(0.140)	(0.045)
Effect (%)	[−38.4%]	[−5.0%]
N observations	298	742
	Close (≤3 km)	Far (>3 km)
Panel B: Distance from City Center (Full Day)		
Night Alert Hours	−0.207***	−0.036
	(0.076)	(0.060)
Effect (%)	[−18.7%]	[−3.5%]
N accidents	750	1016

Notes: Poisson PPML estimates with day-of-week, month, and year fixed effects. Sample: Lviv, February 2022–December 2024. Dependent variable: daily count of traffic accidents with injuries. Panel A compares effects using separate sample splits for weekends (more discretionary travel) vs. weekdays (more commuting). Panel B splits accidents by distance from city center (Rynok Square); central areas have more discretionary destinations. Effect (%) = $\exp(\beta) - 1$. Heteroskedasticity-robust standard errors in parentheses. * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

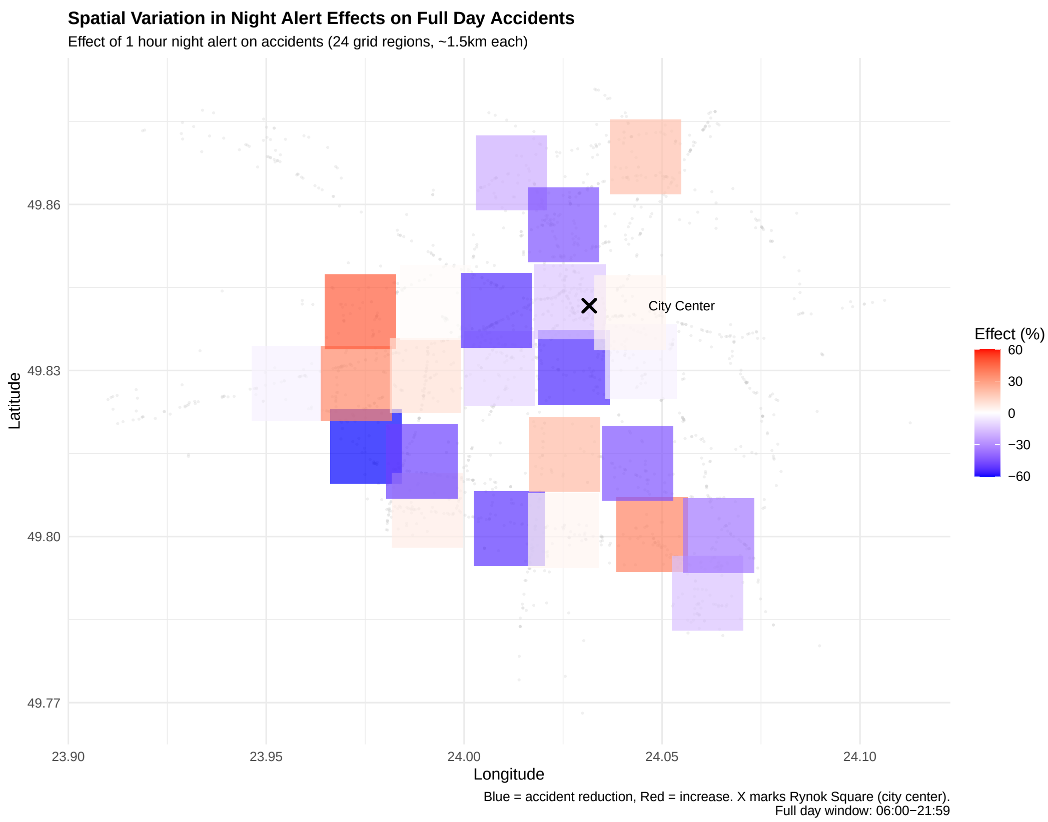

Figure 3: Spatial Heterogeneity in Night Alert Effects

Notes: Each cell shows the estimated effect of night alert hours on full-day accidents in that approximately 1.5 km × 1.3 km grid cell. Blue indicates protective effects (accident reduction); red indicates null or adverse effects. Estimates are from separate Poisson regressions with day-of-week, month, and year fixed effects. Cells with fewer than 30 accidents are excluded.

5.2 Intensive Margin: Cause Composition

The intensive margin predicts that drivers who still travel do so more cautiously. If so, the decline in accidents should fall disproportionately on those caused by risky behaviour. I examine this using the primary-cause codes recorded by the Lviv police.

Each accident record in the police database contains a standardized text description of the primary cause as determined by the investigating officer. I classify accidents into the following cause categories: *fatigue-sensitive* causes (falling asleep at the wheel, driver fatigue, and following too closely due to slow reaction time); *risky behaviour* causes (exceeding safe or posted speed limits, entering oncoming traffic lanes, and improper overtaking); *DUI* (driving under the influence of alcohol); *pedestrian-fault* causes (jaywalking, unexpected entry onto the roadway); *manoeuvring errors* (violations of manoeuvring rules); and *intersection violations* (failure to yield or improper intersection behaviour).[6]

If the drivers become more risk-avoidant, the accidents associated with risky driving ought to fall, while the fatigue-sensitive category, which a sleep-loss channel predicts would rise, should not. Table 4 shows that this is the case. Risky behaviour accidents (speeding, unsafe overtaking) decline by 21.2 percent, while fatigue-sensitive accidents (falling asleep, following too closely due to slow reaction) remain unchanged. Manoeuvring errors decline by 15.6 percent. DUI-related accidents show a large but imprecisely estimated decline (–45.8 percent), possibly because people who would otherwise drink and drive choose not to drink during nights with air alerts or exercise risk avoidance by not driving if they are drunk.

This pattern supports the behavioural risk avoidance hypothesis: people who do drive following night alerts are more careful, avoiding risky manoeuvres, while fatigue-impaired accidents do not increase. An important caveat is that this exercise does not fully distinguish between the selection effect (more cautious drivers continuing to travel while risk-prone drivers stay home) and the change in behaviour among all drivers. The cause-composition

[6]It is important to note that some accidents (approximately 12 percent) are coded into more than one cause category, while for roughly 6 percent the cause was not able to be coded.

evidence is consistent with both interpretations, and I cannot fully distinguish them without individual-level driving data.

Table 4: Effect by Accident Cause Category

	(1) All	(2) Fatigue-sensitive	(3) Risky Behavior	(4) DUI	(5) Pedestrian-Fault	(6) Maneuvering Errors	(7) Intersection Violations
Night Alert Hours	−0.104**	−0.005	−0.239*	−0.613*	−0.193	−0.169*	−0.069
	(0.045)	(0.127)	(0.141)	(0.344)	(0.187)	(0.088)	(0.155)
Effect (%)	[−9.9%]	[−0.5%]	[−21.2%]	[−45.8%]	[−17.6%]	[−15.6%]	[−6.7%]
N accidents	1766	180	248	66	137	691	155

Notes: Poisson PPML estimates for Full Day accidents (06:00–21:59) with day-of-week, month, and year fixed effects. Each column is a separate regression with the indicated accident cause category as the dependent variable. Cause categories are based on official accident reports; the columns shown are the behaviorally interpretable categories and do not exhaust all reported causes, so they do not sum to the All column. "Fatigue-sensitive" includes causes related to driver fatigue or falling asleep. "Risky behavior" includes speeding, unsafe overtaking, and aggressive driving. Effect (%) = $\exp(\beta) - 1$. Heteroskedasticity-robust standard errors in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

5.3 Intensive Margin: Conditional Risk

A complementary approach to the intensive margin uses parking violations as an imperfect but informative proxy for traffic exposure. Parking fines measure a stock of illegally parked vehicles rather than a flow of vehicle-kilometres travelled. However, drivers may consolidate errands, shift trip timing, or reduce mid-day movements while maintaining similar parking patterns, which would attenuate any post-alert changes in the parking data. It is also important to point out that the parking fine series begins only in 2024, so it does not cover the full 2022–2024 sample. The evidence below is therefore suggestive corroboration of the cause-composition results.

Table 5 Panel A compares effects on parking fines versus accidents. Parking fine counts show no significant change following night alerts (Column 1). The full-sample accident estimate shows a significant 9.9 percent decline (Column 2), while the estimate that uses only 2024 data for consistency with the available parking data shows a 9.4 percent decline

(Column 3), directionally consistent but statistically insignificant given the smaller sample. This divergence between parking fines and accidents suggests that traffic volume is roughly unchanged, while accident risk conditional on volume falls.

Panel B formalizes this test using the accident-per-fine ratio as a rough measure of accident risk conditional on traffic activity. Because parking enforcement concentrates in the city centre, the ratio is constructed from central accidents and central fines. The dependent variable is the number of accidents per 100 parking fines (mean ≈ 0.20); the OLS percentage effect is computed as the coefficient divided by the mean of the dependent variable ($-0.051/0.20 \approx -25.4\%$). This estimate is marginally significant at the 10 percent level. A Poisson model with log parking fines as an offset yields a –18.8 percent effect but is imprecisely estimated. While precision is limited given the 2024 parking data sample, both estimates point in the direction of declining risk-taking conditional on driving. A binary any-alert specification of the same Poisson-offset model yields a similar decline of about 21 percent, likewise imprecisely estimated (Table A12).

Another potential concern is that enforcement patterns could change on post-alert days. For instance, police officers may be fatigued from disrupted sleep and reduce their morning patrols, thus leading to an artificial reduction in traffic accident reports and parking ticket reports. I offer several pieces of evidence against this hypothesis. Panel C of Table 5 shows that the effect of night-time air alerts in the morning and afternoon windows is practically identical (and both statistically insignificant).[7] When it comes to the traffic accident reports, the change in the enforcement patterns after the air alert is unlikely. As noted previously, the accident data includes only the cases with physical injuries, and not cases with only property damage that would be easier to dismiss.

[7]Panel C uses the combined 2024–2025 parking fine dataset for greater statistical power, whereas Panels A and B and the appendix parking table (Table A4) use 2024 data only. The time-window breakdown in Appendix Table A4 Panel A (2024 only) shows morning fines unchanged. Thus, the qualitative conclusion (that enforcement timing is stable) holds in both samples.

Table 5: Intensive Margin: Conditional Risk per Unit of Traffic Activity

	(1) Parking Fines (2024)	(2) Accidents (Full Sample)	(3) Accidents (2024)
Panel A: Parking Fines as Activity Proxy (Full Day)			
Night Alert Hours	0.013	−0.104**	−0.099
	(0.040)	(0.045)	(0.111)
Effect (%)	[+1.3%]	[−9.9%]	[−9.4%]
Observations	366	1,040	366
	OLS	Poisson Offset	
Panel B: Accident-per-Fine Ratio			
Night Alert Hours	−0.051*	−0.208	
	(0.028)	(0.154)	
Effect (%)	[−25.4%]	[−18.8%]	
Observations	366		
	(1) Dep. Var.: Morning Fines (06:00–11:59)	(2) Dep. Var.: Afternoon Fines (12:00–17:59)	
Panel C: Enforcement Timing Rule-Out			
Night Alert Hours	−0.021	−0.021	
	(0.013)	(0.014)	
Effect (%)	[−2.1%]	[−2.1%]	
Morning share change	-0.3 pp (p = 0.49)		
Observations	700		
Mean fines/day	589		
Mean accidents/day	1.70		

Notes: Panel A compares the effect of night alerts on parking fine counts (traffic activity proxy, Column 1) vs. accident counts using the full 2022–2024 sample (Column 2) and the 2024-only subsample (Column 3). Column 3 provides the apples-to-apples comparison with parking fines, which are available only for 2024. Panel B tests whether accident risk *conditional on activity* declines, using parking fines as the denominator; because enforcement concentrates in the city center, the ratio is built from central accidents and central fines. The OLS dependent variable is central accidents per 100 central fines (mean ≈ 0.20); the OLS effect is computed as coefficient ÷ mean × 100. The Poisson offset model uses log(central fines) as an offset. Panel C checks whether enforcement timing shifts systematically on alert days; Panel C uses the combined 2024–2025 enforcement dataset for additional coverage. Heteroskedasticity-robust standard errors in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

5.4 Summary of Mechanism Evidence

Taken together, the evidence supports the behavioural risk avoidance interpretation on both margins. On the extensive margin, the concentration of effects on weekends, in the city centre, and in afternoon and evening hours is consistent with residents reducing discretionary travel following night alerts. On the intensive margin, the decline in risky-behaviour accidents coupled with unchanged fatigue-related accidents, and the suggestive evidence from the parking-fine analysis, point to those who do drive exercising greater caution. The concentration of the effect on the day after the alert, with no sign of a persistent or growing response, is more consistent with trip postponement than with a lasting behavioural shift. Appendix Table A12 reports each of these estimates under both the continuous and the binary treatment: the effects tied to the discretionary trip reduction survive under the binary treatment, while those tied to the intensive margin do not, suggesting that the duration of the alert may matter more for the cautious driving than for the trip cancellation.

Overall, the results also show no support for the fatigue-related increase in the accidents, unlike other literature (Smith, 2016; Flynn et al., 2025). I provide additional evidence against fatigue-related accident increase in Table A17. I decompose night exposure into an early window (22:00–02:00), when residents are more likely to be awake to experience the alert, and a late window (02:00–06:00), when alerts are more likely to interrupt deep sleep. A sleep-disruption channel predicts that the late window should matter more, since those are the costliest hours to lose. The data point the other way: early-night exposure is associated with an 18 percent decline in next-day accidents per hour, while late-night exposure shows a smaller and insignificant 8 percent decline. The early-late difference is imprecisely estimated, so I read this as corroboration rather than an independent test, but the pattern again favours the salience of the threat over sleep loss as the main channel.

6 Generalization and Adaptation

The Lviv results establish a protective response and trace it to risk avoidance, but they leave two questions: whether the effect generalizes beyond a single city, and what bounds it. This section answers both through the lens of adaptation. The protective response is not a fixed property of air alerts; it is strongest when alerts are novel and salient and fades as they become routine. I document this fading at three timescales: across cities that differ in how often alerts sound, over time within Lviv, and across consecutive alert nights within a single spell.

6.1 Generalization Across Cities

The cross-city panel, introduced in Section 2.4, constructs the outcome as the count of all reported accidents and the treatment as night alert exposure defined in the same way as for Lviv. Unlike the Lviv injury-only records, this all-accident outcome also includes property-damage collisions. Because these records support a reduced-form comparison but not the Lviv-style mechanism decomposition, and because of this difference in coverage, the cross-city outcome is not comparable in level to the Lviv estimates. However, for comparability purposes, I re-estimate the cross-city regressions, restricting the sample to the injury-and-fatal subset as a robustness check (Section 7). Table 6 summarizes the thirteen cities, which range widely in size, alert frequency, and strike risk.

Table 7 reports pooled PPML estimates with city, day-of-week, month, and year fixed effects. Panel A pools all thirteen cities. Using a binary indicator for any night alert, full-day accidents fall by 3.8 percent and morning accidents by 7.3 percent, both significant with standard errors clustered by city; the afternoon and evening windows are negative but not significant. The continuous treatment yields much smaller and insignificant pooled coefficients. This null, however, is driven by the cities with very frequent air alerts. When the sample is restricted to five low-exposure western cities (Panel B), the results are consistent

Table 6: Cross-City Descriptive Statistics

	Days	Total accidents	Accidents per day	Alert nights (%)	P(strike | alert)
Odesa	1,406	26,607	16.9	56.6	24.4%
Dnipro	1,406	23,934	15.3	95.7	34.5%
Zaporizhzhia	1,406	13,537	8.8	89.3	41.0%
Vinnytsia	1,406	6,625	4.2	41.6	2.5%
Mykolaiv	1,406	6,152	3.9	78.3	19.5%
Chernivtsi	1,406	5,583	3.1	17.6	0.0%
Rivne	1,406	5,480	3.5	21.7	1.2%
Ternopil	1,406	5,153	3.2	20.1	0.8%
Lutsk	1,406	5,044	3.2	21.9	1.1%
Poltava	1,406	4,709	3.1	76.4	1.8%
Sumy	1,406	4,236	2.7	76.3	14.8%
Chernihiv	1,406	3,488	2.2	57.1	11.7%
Uzhhorod	1,406	3,373	2.0	12.4	0.6%

Notes: The thirteen extension cities, sorted by total accident count. "Total accidents" and "Accidents per day" count all reported accidents over the sample period (Feb 2022–Dec 2025). "Alert nights" is the share of nights with any air alert; "P(strike | alert)" is the share of alert days coinciding with a strike within 25 km (VIINA), counting airstrikes and UAVs for all cities plus artillery for the six frontline cities (Dnipro, Mykolaiv, Sumy, Zaporizhzhia, Chernihiv, Odesa). Lviv, the high-detail core city, is summarized separately in Table 1; its accident counts are injury-and-fatal only and so are not comparable in level to the all-accident counts shown here.

with Lviv, and the next-day response increases with the hours of disruption. One difference between cross-city results and Lviv results that is worth pointing out is the morning window. While in Lviv, the night-time alerts have no significant effect on the morning accidents, in the cross-city sample, it is the largest effect. There are several potential explanations. First, the cross-city outcome includes property-damage collisions, which Lviv's injury-only records cannot register; restricted to the injury-and-fatal subset, the pooled morning effect is a null just as in Lviv (Section 7.4). Second, Lviv traffic flows may exhibit distinct cross-day patterns due to it being a popular tourist destination: there could be much more travel in the afternoons and evenings to cultural and entertainment events, which are less common in other cities.

Beyond this difference in the morning window, the broader cross-city pattern is the contrast in the size of the protective effect itself, present where alerts are rare and absent where they are frequent. This discrepancy between western and non-western cities could be explained by adaptation. As western cities experience alerts much less frequently (Table 6), the alerts are still relatively novel and salient, whereas in the highly-exposed to alerts non-western cities, the alerts sound on between 42 and 96 percent of nights, making them lose salience quicker. I test this hypothesis in Appendix Table A15, where I interact the alert exposure and a city's alert frequency. The predicted full-day effect is about −2.4 percent in a city where alerts occur on a fifth of nights but essentially zero where they occur on nine nights in ten, and the interaction is statistically significant. Figure 4 shows the same gradient city by city, plotting each city's full-day effect against its alert frequency.

A competing explanation could be the high rate of false alarms: if in non-western cities the alert is less frequently followed by the actual strike, people may stop relying on the alert as a reliable signal of danger and thus stop adjusting their behaviour. However, the data shows that the opposite is true. These frontline cities face the highest realized strike risk in the sample, with alerts coinciding with a strike within 25 km of the city centre on roughly 12 to 41 percent of nights, against under 3 percent in the western cities (Table 6).

Table 7: Cross-City Pooled Estimates: Night Alerts and Next-Day Accidents

	(1) Morning (06:00–11:59)	(2) Afternoon (12:00–17:59)	(3) Evening (18:00–21:59)	(4) Full Day (06:00–21:59)
Panel A: All cities (13 cities, 18,278 city-days)				
Any night alert	−0.076***	−0.018	−0.030	−0.038**
	(0.023)	(0.017)	(0.023)	(0.019)
Effect (%)	[−7.3%]	[−1.8%]	[−3.0%]	[−3.8%]
Night alert hours	−0.007	−0.001	−0.004	−0.004
	(0.009)	(0.005)	(0.011)	(0.007)
Effect (%)	[−0.7%]	[−0.1%]	[−0.4%]	[−0.3%]
Panel B: Low-exposure cities (5 cities, 7,030 city-days)				
Any night alert	−0.092*	−0.002	−0.061***	−0.039***
	(0.048)	(0.029)	(0.018)	(0.013)
Effect (%)	[−8.8%]	[−0.2%]	[−6.0%]	[−3.8%]
Night alert hours	−0.055***	0.001	−0.030**	−0.020***
	(0.018)	(0.008)	(0.013)	(0.005)
Effect (%)	[−5.3%]	[+0.1%]	[−2.9%]	[−2.0%]
City fixed effects	Yes	Yes	Yes	Yes
Day-of-week / Month / Year FE	Yes	Yes	Yes	Yes

Notes: Poisson pseudo-maximum likelihood (PPML) estimates pooled across cities, with city, day-of-week, month, and year fixed effects. The dependent variable is the daily count of all reported traffic accidents in each time window. "Any night alert" is an indicator for any air alert during the preceding night window (22:00–05:59); "Night alert hours" is the total alert duration in hours. Standard errors clustered by city in parentheses; Effect (%) = $\exp(\beta) - 1$, the implied percentage change in expected accidents. Panel B restricts to the five low-exposure western cities (Chernivtsi, Lutsk, Rivne, Ternopil, Uzhhorod). * p<0.10, ** p<0.05, *** p<0.01.

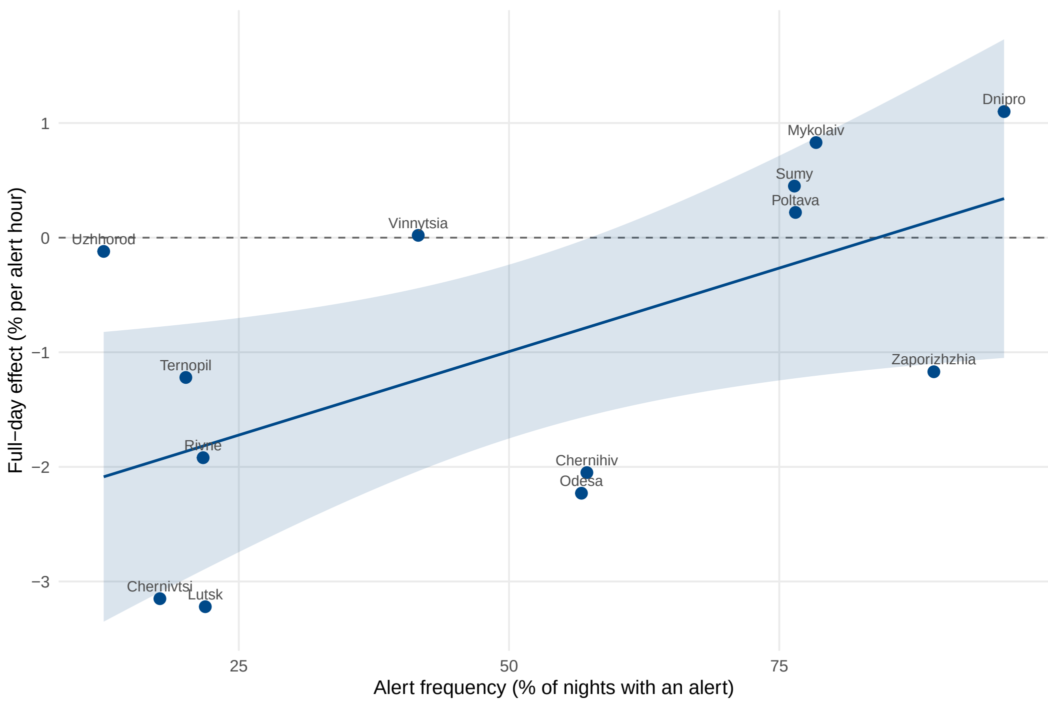

Figure 4: Adaptation Gradient: Full-Day Effect versus Alert Frequency Across Cities

Notes: Each point is one city's full-day effect of night alert exposure (percent change in expected accidents per alert hour) from a city-specific PPML regression with day-of-week, month, and year fixed effects, plotted against the share of nights with any alert. The line is an OLS fit with a 95 percent confidence band (slope = 0.028 percentage points per percentage point of alert frequency, $p = 0.043$, $R^2 = 0.32$). Per-city estimates are reported in Appendix Table A16. Lviv (injury-and-fatal outcome, 17 percent alert nights) is excluded because its outcome is not comparable in level.

6.2 Adaptation Over Time Within Lviv

Adaptation can also be spotted in the Lviv sample. Splitting the Lviv sample by year, the risk-avoidant effect was largest in 2022, the first year of the full-scale invasion, when alerts were novel and closely tied to acute danger: the full-day reduction was about 27 percent, with a particularly large evening decline. In 2023 and 2024, the full-day coefficients remain negative but are statistically insignificant (Table 8). As residents accumulated experience, the marginal alert became less salient, and the protective response weakened, a pattern that parallels Van Dijcke et al. (2023), who document declining civilian responsiveness to Ukrainian air alerts over time. On its own, this within-Lviv attenuation could reflect adaptation or something specific to the first year of the war. The cross-city gradient helps distinguish these interpretations: the effect is weak in heavily alerted cities throughout the sample, so it is familiarity with alerts, not the passage of calendar time, that erodes the response.

6.3 Adaptation Across Consecutive Nights

Finally, adaptation can also be found in the instances where alerts happened on several nights in a row. I classify each alert night by its position within the consecutive set it belongs to: the first night of a run, the second, or the third or later. The protective response is present at full strength from the first night of a run, with no build-up over consecutive nights, and the cumulative alert load over the preceding weeks adds nothing once the immediate night is accounted for (Table A11, Panels A and B respectively). A mechanical-fatigue account predicts the opposite, with accumulated sleep debt deepening the effect over a run of disrupted nights. Instead, the immediate, salient alert matters most, while the cumulative exposure does not bring any additional effect, which points to a behavioural response rather than cumulative exhaustion.

Across all three timescales the pattern is the same: the protective behaviour is a response to a salient threat, subject to habituation, rather than an artifact of the first year of the war or a byproduct of cumulative exhaustion.

Table 8: Year-Specific Treatment Effects (Sample Splits)

	(1) 2022	(2) 2023	(3) 2024
Morning (06–11)			
Night Alert Hours	0.035	0.101	0.055
	(0.163)	(0.113)	(0.177)
Effect (%)	[+3.5%]	[+10.6%]	[+5.7%]
Afternoon (12–17)			
Night Alert Hours	−0.337*	−0.118	−0.101
	(0.178)	(0.091)	(0.152)
Effect (%)	[−28.6%]	[−11.2%]	[−9.6%]
Evening (18–21)			
Night Alert Hours	−0.650**	−0.048	−0.326
	(0.278)	(0.088)	(0.239)
Effect (%)	[−47.8%]	[−4.6%]	[−27.8%]
Full Day (06–21)			
Night Alert Hours	−0.308**	−0.034	−0.099
	(0.123)	(0.049)	(0.111)
Effect (%)	[−26.5%]	[−3.3%]	[−9.4%]
Observations	309	365	366

Notes: Separate PPML regressions estimated for each year subsample. Heteroskedasticity-robust standard errors in parentheses. All specifications include day-of-week and month fixed effects. * p<0.10, ** p<0.05, *** p<0.01.

7 Robustness

7.1 Robustness to Controls

Table 9 gradually adds controls to the baseline specification. Column (1) reproduces the baseline estimate (−9.9 percent). Column (2) addresses a potential concern from the balance tests (Figure A2): five covariates are correlated with night-time alert duration at the 5 percent level (daylight hours, temperature, Lviv city attacks, curfew duration, and daytime alert hours). Column (2) controls for all five directly; the estimate changes minimally to −9.3 percent. The figure examines the principal seasonal and conflict-related covariates. Column (3) adds weather controls (wind speed, precipitation); Column (4) adds the full control set, including the remaining weather variables (humidity, cloud cover, rain, and snow), holidays, school sessions, and PM2.5 air quality. Column (5) replaces separate month and year fixed effects with month-by-year interactions. The estimate remains stable throughout, indicating that the main results are not driven by excluded covariates.

Including daytime alert hours in columns (2)–(5) isolates the night-specific effect: night-time alerts may correlate with "waves" of threat that extend into the following day, so the observed accident reduction could partly reflect mobility reductions during daytime alerts rather than the behavioural response to disrupted sleep. The stability of the night alert coefficient indicates that the results capture a night-specific exposure effect. Notably, neither daytime alert hours nor the Lviv attack indicator is individually significant in the full specification, consistent with the effect operating through the night-time alert experience itself rather than through direct threat exposure or daytime mobility restrictions.

7.2 Identification Diagnostics

I present several diagnostics supporting the identifying assumption that night alert timing is uncorrelated with unobserved determinants of next-day accidents, conditional on fixed effects.

Table 9: Robustness: Weather, Additional Controls, and Alternative Fixed Effects

	(1) Baseline	(2) Imbalanced	(3) Weather	(4) All Controls	(5) Month×Year FE
Panel A: Full Day Accidents (06:00–21:59)					
Night Alert Hours	−0.104**	−0.098**	−0.109**	−0.100**	−0.099**
	(0.045)	(0.046)	(0.045)	(0.046)	(0.048)
Effect (%)	[−9.9%]	[−9.3%]	[−10.3%]	[−9.5%]	[−9.4%]
Panel B: Control Variables					
Day-of-week FE	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	—
Year FE	Yes	Yes	Yes	Yes	—
Month×Year FE	No	No	No	No	Yes
Imbalanced covariates:					
Daylight hours	No	Yes	Yes	Yes	Yes
Temperature	No	Yes	Yes	Yes	Yes
Attack indicator	No	Yes	Yes	Yes	Yes
Curfew duration	No	Yes	Yes	Yes	Yes
Daytime alert hours	No	Yes	Yes	Yes	Yes
Weather controls:					
Wind speed	No	No	Yes	Yes	Yes
Precipitation	No	No	Yes	Yes	Yes
Rain, snow, humidity, cloud	No	No	No	Yes	Yes
Additional controls:					
PM2.5	No	No	No	Yes	Yes
Holiday indicator	No	No	No	Yes	Yes
School session	No	No	No	Yes	Yes
Change from baseline	—	+0.6 pp	-0.4 pp	+0.4 pp	+0.5 pp
Observations	1,040	1,040	1,040	1,040	1,040

Notes: Poisson PPML estimates. Sample: Lviv, February 2022–December 2024. Dependent variable: daily count of traffic accidents with injuries (full day, 06:00–21:59). Column (1) is the baseline specification with calendar fixed effects only. Column (2) controls for all five covariates that predict alert hours at the 5% level (see Figure A2): daylight hours and temperature (seasonal), plus Lviv city attacks, curfew duration, and daytime alert hours (conflict-related). Column (3) adds standard weather controls (wind speed, precipitation). Column (4) adds extended weather controls (rain, snow, humidity, and cloud cover), PM2.5 air quality, holidays, and school sessions. Column (5) replaces separate month and year fixed effects with month-by-year interactions (∼36 parameters), allowing seasonal patterns to vary across years, retaining the full control set from column (4). The stability of estimates across specifications indicates that neither covariate imbalance, omitted weather or institutional factors, time-varying seasonality, nor correlated daytime alert activity drive the main results. Heteroskedasticity-robust standard errors in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

Alert predictability. Table A5 regresses a next-night alert indicator on lagged accident counts and observable local characteristics, including weather, PM2.5, holidays, and conflict-related variables. The pseudo-R^2 reaches only 0.10 even with all available local controls, and the only individually significant predictors are daylight hours and curfew duration—both seasonal and conflict-related variables absorbed by the fixed effects in the main specification. The low predictability of alerts by local factors provides additional support for the identification assumption.

Geographic correlation. Figure A1 shows correlations between Lviv alerts and alerts in all other Ukrainian regions. Correlations are highest with neighbouring western regions and decline with distance, exhibiting a clear geographic gradient that confirms alert timing reflects national-scale military dynamics rather than Lviv-specific factors.

Falsification tests. The lead falsification test (Table A6) includes both the contemporaneous treatment (at t) and tomorrow's alert (at $t + 1$) in the same specification. The contemporaneous coefficients are consistent with the main results, while the lead coefficients are small and statistically insignificant across all time windows, confirming no anticipation effects. Table A3 offers additional falsification tests: tomorrow night's alerts do not predict today's accidents, and night-time alerts do not affect contemporaneous night-time accidents—expected given the night-time curfew that restricts movement.

Attacks as treatment. If the accident reduction reflected the direct effects of strikes (damage, road closures, displacement) rather than the behavioural response to alerts, realized attacks should predict next-day accidents at least as well as alerts do. Table A10 shows they do not: entering prior-day or same-day strikes on Lviv as the treatment yields small, statistically insignificant coefficients, whereas the night-alert effect is unchanged. The effect operates through the alert, not through attack damage.

7.3 Additional Specification Checks

Administrative reporting. The decline reflects behaviour, not changes in reporting, such as strained police resources following high-threat nights. Several features of the setting support this. Lviv is located in western Ukraine, far from the front lines, so police operations are not directly disrupted by alert activity. The null effect on morning accidents argues against systematic underreporting: if recording capacity declined uniformly after alert nights, we would expect reductions across all time windows rather than the afternoon-and-evening concentration observed. As discussed in Section 5.3, parking fine issuance is also unchanged on post-alert days, suggesting stable institutional functioning more broadly. The severity composition is also inconsistent with a reporting artifact (Table A8): the decline is statistically significant for light injury accidents, and the severe-injury point estimate is negative and larger in magnitude though imprecise, the opposite of what a relaxation of reporting thresholds on minor accidents would produce.

OLS specification. Appendix Table A9 replicates the main results using OLS instead of Poisson: the full-day effect is nearly identical to the Poisson estimate, with the same pattern of null morning effects and significant afternoon and evening reductions.

Serial correlation. A potential concern with daily time-series data is serial correlation in residuals, which would not bias point estimates but could affect standard errors. Table A7 Panel B shows that Newey-West and block bootstrap standard errors are very similar to the baseline heteroskedasticity-robust standard errors, and the significance of estimates is unchanged. Panel C reports Augmented Dickey-Fuller tests on Pearson residuals, which strongly reject the unit root null for all outcomes, supporting stationarity.

Randomization inference. As a distribution-free check, I re-estimate the main specification under 1,000 permutations of the treatment and compare the observed t-statistic to its permutation distribution; a small randomization p-value means placebo treatments rarely reproduce an effect as large as the observed one, confirming the baseline inference. I use two schemes—shuffling alert hours within year-by-month-by-day-of-week cells, and circular shifts of the whole treatment series (more conservative, as they preserve the serial correlation of alert spells). Under both, the full-day and afternoon effects are significant at the 5 percent level, the evening effect marginally so, and the morning effect a null (Appendix Table A13). The baseline inference is thus not an artifact of distributional assumptions or of serial dependence in alert timing.

7.4 Cross-City Robustness

The checks above concern the Lviv estimates. The pooled cross-city result of Section 6 raises a separate set of concerns, which I address here.

Inference and fixed effects. The pooled estimates already use standard errors clustered by city. Replacing the separate city and year fixed effects with city-by-year interactions, which absorb any city-specific trend, leaves them essentially unchanged: under the binary treatment the morning reduction is about 6.5 percent and the full-day reduction about 3.3 percent, both significant.

Dynamics and pre-trends. Figure A3 reports a local-projections event study pooled across the thirteen cities, the cross-city analogue of the Lviv event study above. Pre-treatment coefficients are flat, the effect appears on the day after a night alert, and it reverts immediately, mirroring the Lviv dynamics and supporting the identifying assumption in the pooled sample.

Falsification. A lead test that enters tomorrow's alert alongside today's is clean under the continuous treatment: tomorrow's alert hours do not predict today's accidents. Under the binary treatment the lead is significant, but only because alert occurrence is serially correlated, so that an alert tomorrow is itself predicted by an alert today; for this reason I rely on the continuous treatment and the event study, both of which condition on the treatment path, for the dynamic identification check.

Injuries and fatalities sample. Lviv's accident records cover injuries and fatalities only, whereas the cross-city data also include property-damage accidents, so the two outcomes are not directly comparable. Restricting the cross-city sample to the injury-and-fatal subset, the pooled protective effect disappears. The pooled injury sample is powered to detect a Lviv-sized effect at conventional levels, and its full-day confidence interval excludes a decline of that magnitude (Appendix Table A14). This suggests that the all-accident decline is concentrated in property-damage collisions. I also examine whether the injuries-and-fatalities effect replicates in the five low-exposure western cities, the only ones comparable to Lviv (alerts on 12 to 22 percent of nights, against 17 percent in Lviv). However, in those cities the number of injuries is quite small, with daily injury counts of 0.06 to 0.54 against Lviv's 1.7, implying a minimum detectable full-day effect of about 15 percent and power of one half against a Lviv-sized effect. A signal emerges only in the morning window: injury accidents in the western cities fall by 15.2 percent, significant under city-clustered standard errors though imprecise under heteroskedasticity-robust ones. Appendix Table A14 reports these calculations.

8 Conclusion

This paper asks whether exposure to a recurring, salient threat makes people drive more or less dangerously, a question on which the evidence on threat and risk-taking has pointed both ways. Using daily administrative data from Lviv and the unpredictable timing of wartime air alerts, I estimate that each additional hour of overnight alert exposure is associated with roughly 10 percent fewer next-day accidents, concentrated in discretionary afternoon and evening travel and absent in the morning commute. The pattern is consistent with behavioural risk avoidance on two margins: fewer discretionary trips, and more cautious driving among those who do travel, with accidents attributed to risky behaviour falling while those attributed to fatigue do not rise.

Extending the analysis to thirteen other Ukrainian cities turns the within-Lviv attenuation from a puzzle into the paper's central qualification. The protective response is not a fixed feature of air alerts; it fades as alerts become routine. This adaptation is visible at three timescales that point the same way: across years within Lviv, the effect was largest in the first year of the invasion; across cities, it is present where alerts are rare and absent where they are frequent; and within an alert spell, it survives consecutive nights but is keyed to the immediate, salient alert rather than to accumulated exposure. The behavioural response is a reaction to a salient, non-routine threat, not a mechanical consequence of cumulative sleep loss.

Taken together, the evidence indicates that a salient threat can make behaviour systematically safer: the risk-avoidance response dominates any fatigue impairment from the disrupted sleep, reversing the sign one would expect from sleep loss alone. How large the response is depends on how salient the alert remains, which erodes with repetition.

Several limitations bound these conclusions. I observe accidents but not travel, so the extensive margin is identified indirectly, through where and when accidents fall rather than from a direct measure of trips, and the intensive-margin evidence cannot fully separate more careful driving from selection into who drives. The cross-city gradient is descriptive, since cities differ in many ways besides how often alerts sound.

The findings nonetheless carry a clear message for public warning systems. Their effect on behaviour, and hence on road safety, depends not on whether alerts are issued but on whether they remain salient. Rare alerts elicit a protective response; alerts that have become constant and routine do not. This protective response is not free: it could result in foregone and postponed travel, so the same behaviour that lowers accident risk carries a cost in mobility and economic activity. Weighing that trade-off, and understanding how warning systems can preserve the salience that makes them behaviourally effective, is a useful direction for future work.

References

Vaibhav Anand. Does getting forecasts earlier matter? Evidence from winter advisories and vehicle crashes. *American Economic Journal: Economic Policy*, 17(4):106–134, 2025. doi: 10.1257/pol.20230247.

Gary S Becker and Yona Rubinstein. Fear and the response to terrorism: An economic analysis. CEP Discussion Paper No. 1079, London School of Economics, 2011.

Christian Bünnings and Valentin Schiele. Spring forward, don't fall back: The effect of daylight saving time on road safety. *The Review of Economics and Statistics*, 103(1): 165–176, 2021. doi: 10.1162/rest_a_00873.

Jana Cahlíková and Lubomír Cingl. Risk preferences under acute stress. *Experimental Economics*, 20(1):209–236, 2017. doi: 10.1007/s10683-016-9482-3.

Lisa Cameron and Manisha Shah. Risk-taking behavior in the wake of natural disasters. *Journal of Human Resources*, 50(2):484–515, 2015. doi: 10.3368/jhr.50.2.484.

Shooshan Danagoulian and Monica Deza. Driving under the influence of allergies: The effect of seasonal pollen on traffic fatalities. *Journal of Health Economics*, 99:102945, 2025. doi: 10.1016/j.jhealeco.2024.102945.

Thomas Dohmen, Hartmut Lehmann, and Norberto Pignatti. Time-varying individual risk attitudes over the Great Recession: A comparison of Germany and Ukraine. *Journal of Comparative Economics*, 44(1):182–200, 2016.

Jeffrey S Ferris and David A Newburn. Wireless alerts for extreme weather and the impact on hazard mitigating behavior. *Journal of Environmental Economics and Management*, 82:239–255, 2017. doi: 10.1016/j.jeem.2016.11.002.

James Flynn, Noah Meyers-Richter, and Peter Nencka. Soccer's record on the road: The effect of late-night sporting events on fatal car crashes. SSRN Working Paper 5392707, 2025.

Ahmet Gulek. Driving while hungry: The effect of fasting on traffic accidents. *Journal of Development Economics*, 168:103260, 2024. doi: 10.1016/j.jdeveco.2024.103260.

Chie Hanaoka, Hitoshi Shigeoka, and Yasutora Watanabe. Do risk preferences change? Evidence from the Great East Japan Earthquake. *American Economic Journal: Applied Economics*, 10(2):298–330, 2018. doi: 10.1257/app.20170048.

Pamela Jakiela and Owen Ozier. The impact of violence on individual risk preferences: Evidence from a natural experiment. *Review of Economics and Statistics*, 101(3):547–559, 2019.

Òscar Jordà and Alan M. Taylor. Local projections. *Journal of Economic Literature*, 63(1): 59–110, 2025. doi: 10.1257/jel.20241521.

Margaryta Klymak, Andrew Kosenko, Oleg Korenok, Dariia Mykhailyshyna, and Kathryn Vasilaky. Charitable giving in wartime: Evidence from Ukraine's war fundraising. Technical report, Centre for the Study of African Economies, University of Oxford, 2025.

Vadym Klymenko. Ukrainian air raid sirens dataset. `https://github.com/Vadimkin/ukrainian-air-raid-sirens-dataset`, 2024. Accessed: 2026-01-08.

Seung-Keun Martinez, Monika Pompeo, Roman M Sheremeta, Volodymyr Vakhitov, Matthias Weber, and Nataliia Zaika. Civilian evacuation during war: Evidence from Ukraine. *The Economic Journal*, page ueaf075, 2025.

José Luis Montiel Olea and Mikkel Plagborg-Møller. Local projection inference is simpler and more robust than you think. *Econometrica*, 89(4):1789–1823, 2021. doi: 10.3982/ECTA18756.

Andrés Moya. Violence, psychological trauma, and risk attitudes: Evidence from victims of violence in Colombia. *Journal of Development Economics*, 131:15–27, 2018.

Whitney K Newey and Kenneth D West. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708, 1987. doi: 10.2307/1913610.

Issi Romem and Ity Shurtz. The accident externality of driving: Evidence from observance of the Jewish Sabbath in Israel. *Journal of Urban Economics*, 96:36–54, 2016. doi: 10.1016/j.jue.2016.07.004.

Yau-Huo Shr, Wen Hsu, Bing-Fang Hwang, and Chau-Ren Jung. Air quality and risky behaviors on roads. *Journal of Environmental Economics and Management*, 118:102786, 2023. doi: 10.1016/j.jeem.2023.102786.

Austin C Smith. Spring forward at your own risk: Daylight saving time and fatal vehicle crashes. *American Economic Journal: Applied Economics*, 8(2):65–91, 2016. doi: 10.1257/app.20140100.

Guy Stecklov and Joshua R Goldstein. Terror attacks influence driving behavior in Israel. *Proceedings of the National Academy of Sciences*, 101(40):14551–14556, 2004. doi: 10.1073/pnas.0402483101.

Guy Stecklov and Joshua R Goldstein. Societal responses to endemic terror: Evidence from driving behavior in Israel. *Social Forces*, 88(4):1859–1884, 2010. doi: 10.1353/sof.2010.0027.

David Van Dijcke, Austin L Wright, and Mark Polyak. Public response to government alerts saves lives during Russian invasion of Ukraine. *Proceedings of the National Academy of Sciences*, 120(18):e2220160120, 2023. doi: 10.1073/pnas.2220160120.

Maarten J Voors, Eleonora E M Nillesen, Philip Verwimp, Erwin H Bulte, Robert Lensink, and Daan P Van Soest. Violent conflict and behavior: a field experiment in Burundi. *American Economic Review*, 102(2):941–964, 2012.

Yuri Zhukov and Natalie Ayers. VIINA 2.0: Violent Incident Information from News Articles on the 2022 Russian Invasion of Ukraine. Cambridge, MA: Harvard University, 2023. `https://github.com/zhukovyuri/VIINA`, accessed 10 February 2026.

Online Appendix (Not for Publication)

Table A1: Distributed Lag Effects: Persistence of Night Alert Impact

	(1) Morning	(2) Afternoon	(3) Evening	(4) Full Day
h=0 (Same day)	0.039	−0.168**	−0.168*	−0.108**
	(0.086)	(0.073)	(0.087)	(0.046)
h=1 (Next day)	−0.090	0.022	−0.039	−0.024
	(0.119)	(0.071)	(0.087)	(0.050)
h=2 (+2 days)	0.032	0.097	0.121	0.092**
	(0.098)	(0.063)	(0.088)	(0.043)
h=3 (+3 days)	−0.064	0.058	0.077	0.037
	(0.094)	(0.063)	(0.071)	(0.041)

Notes: Poisson PPML estimates with day-of-week, month, and year fixed effects. Each column is a single regression including all four lags simultaneously. Effect (%) = $\exp(\beta) - 1$. The effect is concentrated at h=0 (same day) with no persistence at h=1, and a positive reversal at h=2 consistent with catch-up travel. Heteroskedasticity-robust standard errors in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

Table A2: Heterogeneity by Area Type

	(1) City Centre	(2) Transport Hub	(3) Residential
Full Day Accidents (06:00–21:59)			
Night Alert Hours	−0.355*	−0.015	−0.086
	(0.181)	(0.113)	(0.056)
Effect (%)	[−29.9%]	[−1.5%]	[−8.2%]
N accidents	249	320	1196

Notes: Poisson PPML estimates with day-of-week, month, and year fixed effects. Area types are classified based on location characteristics. City Centre (Halytskyi district) includes the historic center and main commercial areas. Transport Hub (Zaliznychnyi district) includes areas near the main railway station. Residential includes the remaining four districts. Effect (%) = $\exp(\beta)-1$. Heteroskedasticity-robust standard errors in parentheses. * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

Table A3: Falsification Tests: Lead Effects and Night Placebo

	(1) Lead Test (Tomorrow → Today)	(2) Night Placebo (Same-Night Effect)
Night Alert Hours	−0.011	−0.053
	(0.045)	(0.161)
Effect (%)	[−1.1%]	[−5.1%]
Dep. Var.	Full Day Accidents	Night Accidents
Result	PASS — Null	Contemporaneous effect

Notes: Poisson PPML estimates with day-of-week, month, and year fixed effects. Column (1) tests whether tomorrow's night alert predicts today's accidents (lead falsification). Under the causal model, lead effects should be zero. The null result supports identification. Column (2) tests the contemporaneous effect on same-night accidents during alert hours. The estimated coefficient is small and statistically insignificant, indicating no detectable change in accident rates during the alert window itself. This null same-night result is distinct from the next-day behavioral response documented in the main results. Heteroskedasticity-robust standard errors in parentheses. * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

Table A4: Traffic Activity Proxy: Parking Fine Analysis

	(1) Morning (06:00–11:59)	(2) Afternoon (12:00–17:59)	(3) Full Day (06:00–21:59)
Panel A: Parking Fines by Time Window			
Night Alert Hours	−0.015	0.028	0.013
	(0.046)	(0.047)	(0.040)
Effect (%)	[−1.5%]	[+2.9%]	[+1.3%]
	(1) Parking Fines (2024)	(2) Accidents (Full Sample)	(3) Accidents (2024)
Panel B: Comparison with Accidents (Full Day)			
Night Alert Hours	0.013	−0.104**	−0.099
	(0.040)	(0.045)	(0.111)
Effect (%)	[+1.3%]	[−9.9%]	[−9.4%]
Observations	366	1,040	366
Mean fines/day		589	

Notes: Poisson PPML estimates with day-of-week, month, and year fixed effects. Parking fines serve as a proxy for traffic activity/exposure. Both panels use 2024 parking data only; the enforcement timing test in Table 5 Panel C uses the extended 2024–2025 combined dataset for additional coverage. Panel B Column (1) uses 2024 parking fines; Column (2) uses the full-sample (2022–2024) accident estimate; Column (3) restricts to 2024 only for an apples-to-apples comparison with parking fines. The 2024-only accident estimate is directionally consistent with the full-sample result but imprecisely estimated due to the smaller sample. Heteroskedasticity-robust standard errors in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

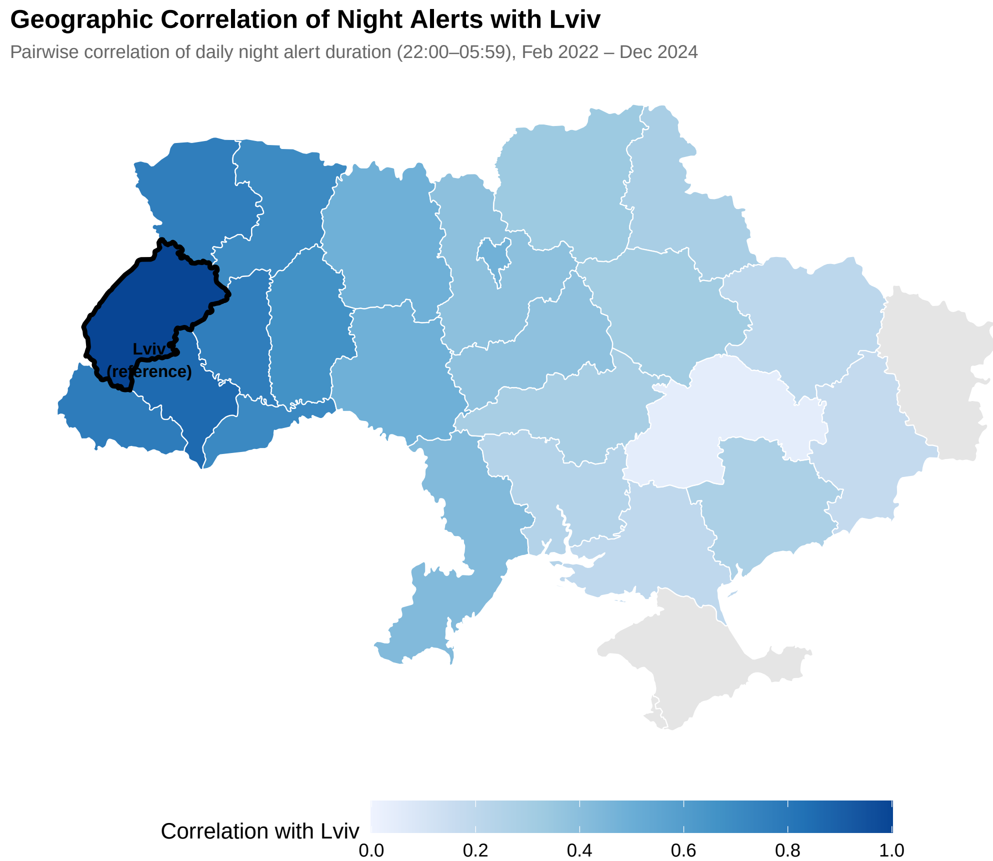

Figure A1: Geographic Correlation of Night Alerts with Lviv

Notes: Each region is shaded by its Pearson correlation of daily night alert duration (22:00–05:59) with Lviv over the sample period (February 2022–December 2024). Lviv (reference, outlined in black) has a correlation of 1.0 by construction. Correlations are highest with neighboring western regions and decline with distance, exhibiting a clear geographic gradient. This pattern confirms that alert timing reflects national-scale military dynamics (large aerial attacks trigger alerts across multiple regions simultaneously) rather than Lviv-specific conditions. Regions with no alert data (Crimea, Sevastopol) are shown in gray. Correlations for partially occupied oblasts (Donetsk, Luhansk, Zaporizhzhia, Kherson) reflect government-controlled territory.

Table A5: Predictability of Night Alerts

	(1) Lagged Accidents	(2) Weather	(3) Holiday	(4) All Controls
Lagged Full-Day Accidents	−0.041	−0.080	−0.043	0.003
	(0.062)	(0.064)	(0.062)	(0.067)
Temperature		0.040***		−0.009
		(0.010)		(0.020)
Precipitation		−0.033		−0.035
		(0.029)		(0.030)
Holiday			−1.068	−1.257*
			(0.740)	(0.750)
Wind Speed				−0.015
				(0.020)
Daylight Hours				0.259***
				(0.062)
Daytime Alert Hours				0.106
				(0.104)
Curfew Duration				0.527***
				(0.117)
Lviv City Attack				0.051
				(0.242)
School Session				0.290
				(0.224)
PM2.5				0.010
				(0.022)
Day-of-week FE	Yes	Yes	Yes	Yes
Pseudo R^2	0.010	0.029	0.013	0.104
Observations	1,038	1,038	1,038	1,038

Notes: Logistic regression. Dependent variable: binary indicator for any night alert (22:00–05:59). All specifications include day-of-week fixed effects. Column (1) adds lagged full-day accident count. Column (2) adds weather controls (temperature, precipitation). Column (3) adds a holiday indicator. Column (4) includes all available local predictors: lagged accidents, weather (temperature, precipitation, wind speed), daylight hours, holiday, daytime alert hours, curfew duration, Lviv city attacks, school session, and PM2.5 air quality. Daylight hours and curfew duration are individually significant in Column (4), but both reflect seasonality and conflict dynamics absorbed by month and year fixed effects in the main outcome regressions. * p<0.10, ** p<0.05, *** p<0.01.

Table A6: Falsification Test: Lead Night Alert Effects

	(1) Morning (06:00–11:59)	(2) Afternoon (12:00–17:59)	(3) Evening (18:00–21:59)	(4) Full Day (06:00–21:59)
Night Alert Hours (at t)	0.033	−0.161**	−0.162*	−0.104**
	(0.085)	(0.073)	(0.088)	(0.045)
Lead Night Alert Hours (at $t+1$)	0.015	−0.007	−0.046	−0.011
	(0.089)	(0.070)	(0.093)	(0.045)
Effect (%), lead	[+1.5%]	[−0.7%]	[−4.5%]	[−1.1%]
Observations	1,040	1,040	1,040	1,040

Notes: Poisson PPML estimates with day-of-week, month, and year fixed effects. Both contemporaneous (at t) and lead (at $t+1$) night alert hours are included in the same specification. Lead Night Alert Hours measures tomorrow night's alert exposure. Under the causal model, lead effects should be zero—tomorrow's alert cannot cause today's accidents. The contemporaneous coefficient is consistent with the main results (Table 2). The null lead results across all time windows support identification. Heteroskedasticity-robust standard errors in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

Table A7: Robustness: Alternative Standard Errors and Stationarity

	(1) Morning (06:00–11:59)	(2) Afternoon (12:00–17:59)	(3) Evening (18:00–21:59)	(4) Full Day (06:00–21:59)
Panel A: Baseline Estimates				
Night Alert Hours	0.034	−0.162**	−0.163*	−0.104**
	(0.085)	(0.073)	(0.087)	(0.045)
Panel B: Alternative Standard Errors				
Newey-West	(0.080)	(0.073)	(0.087)	(0.047)
Block bootstrap	(0.086)	(0.073)	(0.093)	(0.050)
Panel C: Stationarity (Augmented Dickey-Fuller)				
ADF statistic	-10.25	-9.65	-9.91	-9.71
ADF p-value	<0.01	<0.01	<0.01	<0.01
Observations	1,040	1,040	1,040	1,040

Notes: Panel A reports the baseline Poisson PPML coefficient and its heteroskedasticity-robust standard error. Panel B reports alternative standard errors for that coefficient: Newey-West standard errors use the Newey and West (1987) plug-in bandwidth to account for heteroskedasticity and autocorrelation, and block bootstrap standard errors resample weekly blocks (7 consecutive days) with replacement (999 replications). Panel C reports Augmented Dickey-Fuller tests on Pearson residuals from the baseline model; H_0: unit root (non-stationary). Rejection at all conventional levels confirms that residual series are stationary. All specifications include day-of-week, month, and year fixed effects. * p<0.10, ** p<0.05, *** p<0.01.

Table A8: Effect of Night Alert Hours on Accidents by Severity Category

	(1) All	(2) Light	(3) Severe	(4) Fatal
Night Alert Hours	−0.104**	−0.100**	−0.318	0.161
	(0.045)	(0.049)	(0.220)	(0.215)
Effect (%)	[−9.9%]	[−9.5%]	[−27.2%]	[+17.5%]
Observations	1,040	1,040	1,040	1,040
Day-of-week FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes

Notes: Each column reports a separate Poisson PPML regression where the dependent variable is the daily count of accidents in the indicated severity category. Night Alert Hours measures total hours of alert exposure during 22:00–05:59. "Light" = non-severe injuries. "Severe" = severe non-fatal injuries. Heteroskedasticity-robust standard errors in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

Table A9: OLS Robustness: Main Results

	(1) Morning (06–12)	(2) Afternoon (12–18)	(3) Evening (18–22)	(4) Full Day (06–22)
Night Alert Hours	0.014	−0.096**	−0.068**	−0.150**
	(0.038)	(0.038)	(0.034)	(0.059)
Effect (% of mean)	[+3.3%]	[−12.7%]	[−13.2%]	[−8.8%]
Mean accidents/day	0.43	0.75	0.52	1.70
Observations	1,040	1,040	1,040	1,040
FE	Day-of-week + Month + Year			

Notes: OLS estimates for comparison with the Poisson PPML results in Table 2. Dependent variable: daily count of traffic accidents with injuries. Effect (% of mean) computed as coefficient divided by the outcome mean. Heteroskedasticity-robust standard errors in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

Table A10: Night Alerts versus Realized Attacks as Treatment (Lviv)

	Morning	Afternoon	Evening	Full Day
Night alert hours	0.034	−0.162**	−0.163*	−0.104**
	(0.085)	(0.073)	(0.087)	(0.045)
Effect (%)	[+3.4%]	[−14.9%]	[−15.1%]	[−9.9%]
Prior-day attack	−0.073	0.100	−0.144	−0.010
	(0.139)	(0.094)	(0.142)	(0.065)
Effect (%)	[−7.0%]	[+10.5%]	[−13.4%]	[−1.0%]
Same-day attack	−0.139	−0.112	0.110	−0.051
	(0.145)	(0.103)	(0.117)	(0.066)
Effect (%)	[−12.9%]	[−10.6%]	[+11.7%]	[−5.0%]
Observations	1,039	1,039	1,039	1,039
Day-of-week FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes

Notes: PPML estimates for Lviv. Each panel reports a separate regression of next-day accident counts in the indicated window on the listed treatment: night alert hours (baseline), or an indicator for a realized attack on Lviv (VIINA strikes) on the prior day or the same day. Heteroskedasticity-robust standard errors in parentheses; the bracketed row converts the coefficient to a percentage change, $\exp(\beta) - 1$. Realized attacks do not predict next-day accidents, indicating that the effect operates through the alert (sleep disruption and threat salience) rather than through direct attack damage. * p<0.10, ** p<0.05, *** p<0.01.

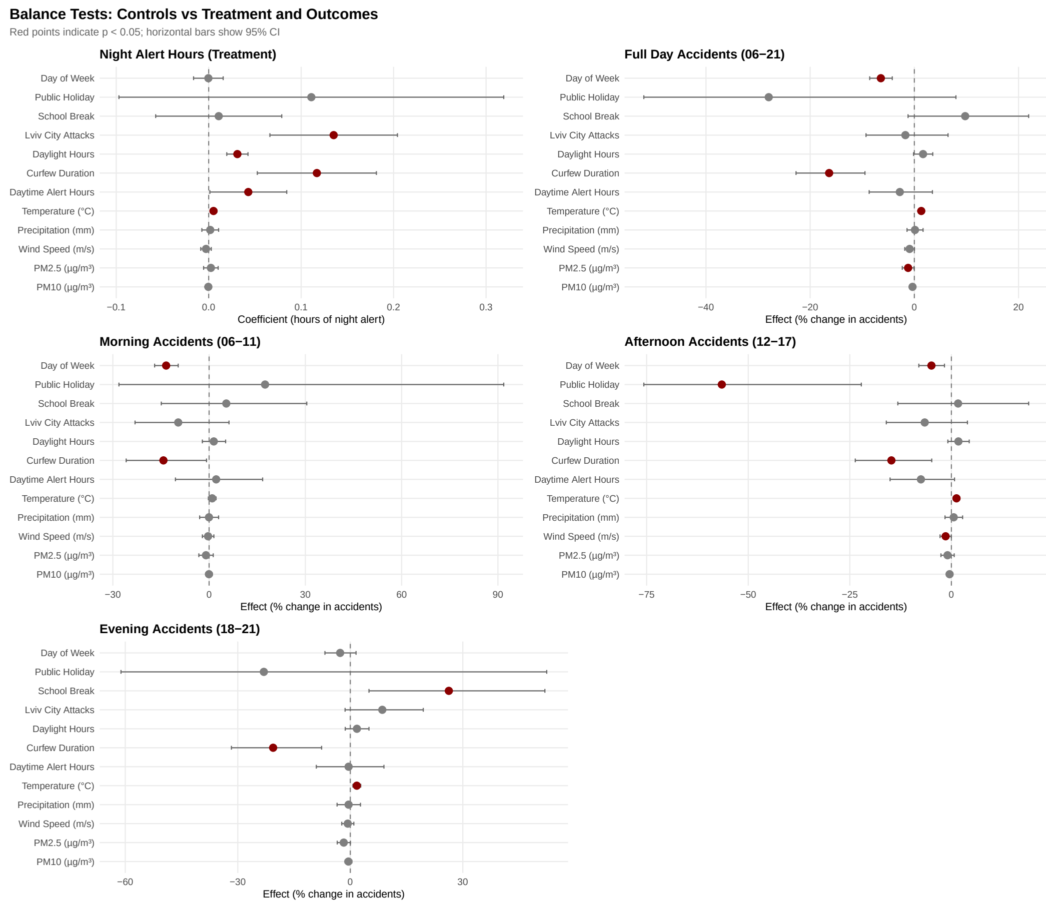

Figure A2: Covariate Balance: Alert vs. Non-Alert Days

Notes: Coefficient estimates and 95 percent confidence intervals from separate univariate regressions of each covariate on night alert hours (OLS for treatment panel; Poisson for outcome panels). Red indicates $p < 0.05$. Daylight hours and temperature show seasonal imbalance (alerts more common in summer), absorbed by month fixed effects. Lviv city attacks, curfew duration, and daytime alert hours reflect shared conflict dynamics. PM2.5 air quality, wind speed, and precipitation are well balanced against the continuous treatment.

Table A11: Position Within a Consecutive Alert Run and Cumulative Alert Load (Cross-City)

	Full Day	Morning
Panel A: Position within a consecutive alert run		
First night of run	−0.043**	−0.087***
	(0.017)	(0.025)
Effect (%)	[−4.2%]	[−8.4%]
Second night of run	−0.050***	−0.072***
	(0.014)	(0.027)
Effect (%)	[−4.8%]	[−7.0%]
Third or later night	−0.029	−0.068***
	(0.024)	(0.025)
Effect (%)	[−2.9%]	[−6.6%]
Panel B: Recent cumulative load (prior 30-day alert count)		
Tonight's alert (any)	−0.036***	−0.071***
	(0.008)	(0.013)
Effect (%)	[−3.5%]	[−6.9%]
Prior 30-day alert count	−0.001	−0.002
	(0.005)	(0.006)
Effect (%)	[−0.1%]	[−0.2%]

Notes: Pooled PPML estimates across the thirteen extension cities (city, day-of-week, month, and year fixed effects); the dependent variable is the daily count of all reported accidents. Standard errors clustered by city in parentheses; Effect (%) = $\exp(\beta) - 1$. *Panel A* classifies each alert night by its position within the consecutive-alert run it belongs to: the first night of a run (no alert the previous night—the opening night of a run, not necessarily an isolated single night), the second consecutive night, or the third or later night; each coefficient is relative to a non-alert night, and the effect is present from the first night of a run and shows no significant attenuation across subsequent consecutive nights. *Panel B* adds the count of alert nights over the preceding 30 days alongside tonight's alert: the cumulative-load coefficient is small and insignificant while tonight's alert retains its effect, so the response is keyed to the immediate, salient night rather than to accumulated exposure; 7- and 14-day windows give the same pattern. * p<0.10, ** p<0.05, *** p<0.01.

Table A12: Continuous versus Binary Treatment Across Specifications (Lviv)

	Continuous (per hour)	Binary (any alert)
Main results, by window		
Morning (06–11)	0.034	0.052
	(0.085)	(0.132)
Effect (%)	[+3.4%]	[+5.3%]
Afternoon (12–17)	−0.162**	−0.094
	(0.073)	(0.096)
Effect (%)	[−14.9%]	[−9.0%]
Evening (18–21)	−0.163*	−0.247*
	(0.087)	(0.129)
Effect (%)	[−15.1%]	[−21.9%]
Full day (06–21)	−0.104**	−0.097
	(0.045)	(0.065)
Effect (%)	[−9.9%]	[−9.3%]
Extensive margin		
Weekend	−0.485***	−0.312**
	(0.140)	(0.159)
Effect (%)	[−38.4%]	[−26.8%]
Weekday	−0.051	−0.057
	(0.045)	(0.073)
Effect (%)	[−5.0%]	[−5.5%]
Center (≤3 km)	−0.207***	−0.221**
	(0.076)	(0.101)
Effect (%)	[−18.7%]	[−19.9%]
Far (>3 km)	−0.036	−0.007
	(0.060)	(0.086)
Effect (%)	[−3.5%]	[−0.7%]
Intensive margin: accident causes		
Risky behavior	−0.239*	0.018
	(0.141)	(0.181)
Effect (%)	[−21.2%]	[+1.8%]
Fatigue-sensitive	−0.005	0.093
	(0.128)	(0.191)
Effect (%)	[−0.5%]	[+9.7%]
DUI (placebo)	−0.613*	−0.209
	(0.344)	(0.360)
Effect (%)	[−45.8%]	[−18.8%]
Intensive margin: conditional risk		
Accident-per-fine ratio	−0.208	−0.239
	(0.154)	(0.152)
Effect (%)	[−18.8%]	[−21.3%]

Notes: Each block reports the treatment coefficient from a separate PPML regression (Lviv, day-of-week, month, and year fixed effects), with heteroskedasticity-robust standard errors in

Table A13: Randomization Inference for the Baseline Estimates

	(1) Morning	(2) Afternoon	(3) Evening	(4) Full Day
Observed t-statistic	0.40	−2.22	−1.87	−2.30
Analytic p-value (robust)	0.691	0.026	0.062	0.021
RI p-value: within-cell permutation	0.735	0.021	0.032	0.015
RI p-value: circular shift	0.716	0.038	0.075	0.023
Permutations	1,000	1,000	1,000	1,000
Observations	1,040	1,040	1,040	1,040

Notes: Randomization inference for the baseline Poisson specification (Table 2, Panel A). The test statistic is the t-statistic on night alert hours (coefficient divided by heteroskedasticity-robust standard error); the randomization p-value is the share of permutations with $|t|$ at least as large as observed, computed as $(\#\{|t_{perm}| \geq |t_{obs}|\} + 1)/(N + 1)$. Two permutation schemes: *within-cell* shuffles night alert hours across days within year × month × day-of-week cells, respecting the fixed-effect structure; *circular shift* rotates the entire treatment series by a random offset of at least 14 days, preserving the serial correlation of alert spells, and is therefore the more conservative scheme. 1,000 permutation draws per scheme, shared across outcomes.

Table A14: Power of the Harmonized Injury-and-Fatal Outcome in the Cross-City Sample

	(1) Morning	(2) Afternoon	(3) Evening	(4) Full Day
Panel A: All 13 cities (N = 18,278 city-days)				
Estimated effect (%)	−1.3	+4.8	−0.8	+1.4
95% confidence interval	[−9.1, +7.1]	[−2.0, +12.1]	[−8.4, +7.6]	[−2.9, +5.9]
MDE at 80% power (%)	12.4	10.1	12.2	6.4
Power vs. Lviv-sized effect	0.64	0.81	0.66	0.99
Panel B: 5 low-exposure western cities (N = 7,030 city-days)				
Estimated effect (%)	−15.2	+0.8	+8.9	−0.9
95% confidence interval	[−30.9, +4.1]	[−12.7, +16.5]	[−7.2, +28.0]	[−9.8, +9.0]
MDE at 80% power (%)	33.9	22.9	25.8	14.5
Power vs. Lviv-sized effect	0.15	0.26	0.22	0.52

Notes: Power analysis for the harmonized injury-and-fatal outcome under the binary any-alert treatment, based on the pooled PPML estimates with city, day-of-week, month, and year fixed effects. The minimum detectable effect (MDE) at 80% power and 5% size is $\exp(2.8 \times SE) - 1$; power against a Lviv-sized effect is $\Phi(|\beta_{Lviv}|/SE - 1.96)$ with $\beta_{Lviv} = -0.097$, the Lviv binary full-day estimate (−9.3%). Both use heteroskedasticity-robust standard errors: cluster-robust standard errors are unreliable for power calculations with only five clusters in Panel B. Western cities: Chernivtsi, Lutsk, Rivne, Ternopil, Uzhhorod.

Table A15: Adaptation: Night Alert Effects Interacted with City Alert Frequency

	Night alert hours		Any night alert	
	Morning (1)	Full Day (2)	Morning (3)	Full Day (4)
Night alert exposure	−0.028***	−0.013***	−0.083***	−0.041**
	(0.008)	(0.005)	(0.024)	(0.018)
Exposure × alert frequency	0.077***	0.037**	0.110	0.075
	(0.021)	(0.015)	(0.072)	(0.055)
Implied effect at 20% alert nights	−5.0%	−2.4%	−11.1%	−6.3%
Implied effect at 90% alert nights	+0.2%	+0.1%	−4.0%	−1.2%
Observations	18,278	18,278	18,278	18,278

Notes: Pooled PPML estimates across the thirteen extension cities with city, day-of-week, month, and year fixed effects. "Alert frequency" is the city-level share of nights with any alert, centered at the pooled sample mean, so the main effect is the effect at mean alert frequency. The implied effects evaluate the linear combination at an alert frequency of 20 percent (typical of the western cities) and 90 percent (typical of the most heavily alerted cities), expressed as the percentage change in expected accidents. Standard errors clustered by city (thirteen clusters) in parentheses, and significance stars are from this clustered inference. For reference, the heteroskedasticity-robust interaction p-values are <0.001, 0.003, 0.132, and 0.073 for columns (1)–(4). * p<0.10, ** p<0.05, *** p<0.01.

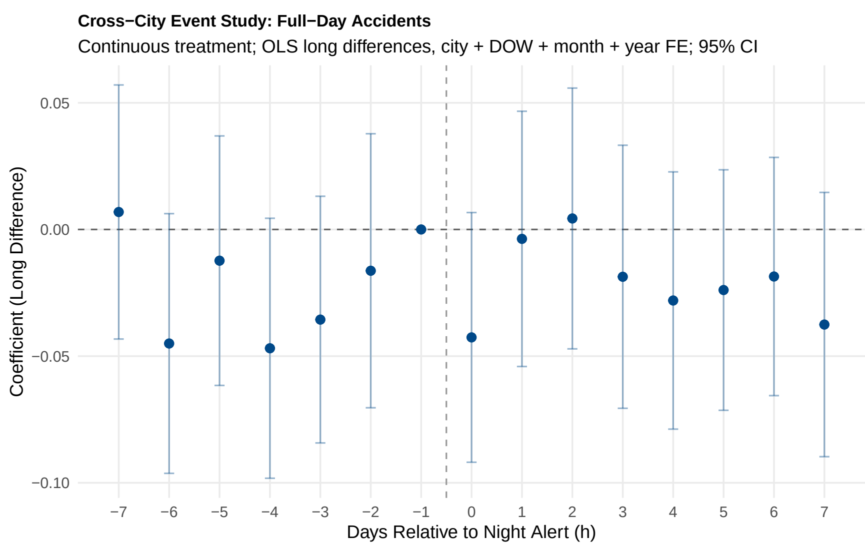

Figure A3: Cross-City Event Study: Dynamic Effect of Night Alerts on Full-Day Accidents

Notes: Local projection estimates pooled across the thirteen extension cities, following Jordà and Taylor (2025), using the continuous (per-hour) night-alert treatment and the full-day outcome. Each point is the long-difference coefficient at horizon h ($y_{t+h} - y_{t-1}$ regressed on the treatment) with city, day-of-week, month, and year fixed effects; bars are 95 percent confidence intervals. Pre-treatment horizons are flat, the effect appears at $h = 0$, and it reverts by $h = +1$, mirroring the Lviv dynamics.

Table A16: Per-City Night Alert Effects, Sorted by Alert Frequency

	Alert nights (%)	Accidents/day	Morning	Full Day
Uzhhorod	12.4	2.0	+7.2% (9.6)	+0.2% (3.7)
Chernivtsi	17.6	3.1	−7.9% (4.9)	−3.2%* (1.9)
Ternopil	20.1	3.2	−1.7% (3.8)	−1.2% (2.3)
Rivne	21.7	3.5	−5.2% (3.4)	−1.9% (2.0)
Lutsk	21.9	3.2	−9.0%** (3.5)	−3.2% (2.0)
Vinnytsia	41.6	4.2	−0.1% (2.1)	+0.0% (1.1)
Odesa	56.6	16.9	−3.0%*** (1.1)	−2.2%*** (0.7)
Chernihiv	57.1	2.2	−1.1% (2.5)	−2.0% (1.5)
Sumy	76.3	2.7	−2.4% (2.2)	+0.5% (1.4)
Poltava	76.4	3.1	−1.3% (2.2)	+0.2% (1.3)
Mykolaiv	78.3	3.9	−0.3% (2.0)	+0.8% (1.1)
Zaporizhzhia	89.3	8.8	−0.3% (1.1)	−1.2% (0.7)
Dnipro	95.7	15.3	+1.6%* (0.9)	+1.2%** (0.5)

Notes: City-by-city PPML estimates of the continuous (per-hour) night alert effect with day-of-week, month, and year fixed effects, expressed as the percentage change in expected accidents per alert hour; delta-method heteroskedasticity-robust standard errors in percentage points in parentheses. Cities are sorted by the share of nights with any alert. "Accidents/day" is the mean number of full-day (06:00–21:59) accidents per day and indexes how thin each city's data are. The outcome is the daily count of all reported accidents; Lviv (injury-and-fatal only, 17.2 percent alert nights) is not shown because its outcome is not comparable in level. Several cities record only two to three accidents per day, so the individual per-city estimates are imprecise and mostly indistinguishable from zero. They are reported as descriptive inputs to the adaptation gradient rather than as standalone results: inference rests on the pooled estimate (Table 7) and the frequency interaction (Table A15), and these per-city points underlie Figure 4. * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

Table A17: Time-of-Night Decomposition: Early versus Late Night Alerts

	(1) Morning	(2) Afternoon	(3) Evening	(4) Full Day
Early-night alert hours (22:00–02:00)	0.048	−0.345**	−0.276	−0.203*
	(0.195)	(0.171)	(0.236)	(0.117)
Effect (%)	[+5.0%]	[−29.2%]	[−24.1%]	[−18.4%]
Late-night alert hours (02:00–06:00)	0.030	−0.116	−0.137	−0.080
	(0.101)	(0.081)	(0.101)	(0.052)
Effect (%)	[+3.1%]	[−10.9%]	[−12.8%]	[−7.7%]
p-value: early = late	0.936	0.236	0.608	0.362
Observations	1,040	1,040	1,040	1,040

Notes: Poisson PPML estimates with day-of-week, month, and year fixed effects. Night alert exposure is decomposed into an early window (22:00 of the preceding evening to 02:00) and a late window (02:00 to 06:00); both exposures enter the same regression, and they sum to total night alert hours. Effect (%) = $\exp(\beta) - 1$ per hour of exposure in the indicated window. The "early = late" row reports the p-value of a Wald test that the two coefficients are equal. Heteroskedasticity-robust standard errors in parentheses. * p<0.10, ** p<0.05, *** p<0.01.