Seleno-Gravitational Rhythm (SGR): Impulsive Lunar Tidal Waves as a Modulator of Human Consciousness, Behavior, and Somatic Health

A Cross-City Meta-Analysis of Criminal and Medical Statistics in U.S. Metropolitan Areas (2001–2025)

Author: Artem Tysiatskii
Status: Independent researcher and software developer
Location: Kaliningrad, Russia
Email: artem.tysiatskii@gmail.com
ORCID: https://orcid.org/0009-0006-1974-7894
GitHub repository: https://github.com/TemaTys/seleno-gravitational-rhythm
Zenodo DOI: `[DOI: 10.5281/zenodo.20518660]`

Version: 1.0
Year: 2026

How to cite this work

Tysiatskii, A. (2026). *Seleno-Gravitational Rhythm (SGR): Impulsive lunar tidal waves as a modulator of human consciousness, behavior, and somatic health*. Zenodo. `[DOI: 10.5281/zenodo.20518660]`

Data and code availability

Source code, processed data tables, wave files, and analytical outputs are maintained in the project repository:
GitHub: https://github.com/TemaTys/seleno-gravitational-rhythm
Zenodo archival snapshot: `[DOI: 10.5281/zenodo.20518660]`

Abstract

The link between the lunar cycle and human physiological and behavioral states has been discussed since the time of ancient medicine, yet the contemporary literature typically restricts itself to the analysis of four static lunar phases (New Moon, Full Moon, quarters) and yields contradictory results — largely because what may be biologically meaningful is not only the absolute level of the tidal background, but also the dynamics of its variation, including the first and second derivatives of the tidal potential. In the present work, we develop and apply, for the first time, an original detector of impulsive jumps in the second derivative of the tidal potential (**ddF**, *lunar tidal-wave detector*) to a large body of official data

— seven independent databases from five U.S. metropolitan areas (Chicago, Los Angeles, New York, Philadelphia, San Francisco) covering the period 2001–2025 and comprising more than 106 million records of criminal statistics, ambulance dispatches, and 911 calls.

Two independent analytical pipelines are employed: a random-effects meta-analysis with HKSJ correction for cross-city replication, and a single-database pipeline equipped with an internal robustness battery (block-permutation, bootstrap, slice stability, placebo) for unique databases. A robust replication signal is obtained (TIER1, $p = 3.9 \times 10^{-3}$, $I^2 = 0\%$, $k = 21$ across 5 cities) for the family of property crimes within the impulsive-wave window preceding the lunar quadrature, together with 14 TIER1 and 17 TIER2 findings in the single-database analysis, reproduced in the 2001–2014 sub-cohort.

A robust **syzygy / quadrature** dichotomy is identified: lunar syzygies and quadratures act on the cardiovascular system and property crime with opposite signs — a structural fingerprint that is difficult to reconcile with random noise. We propose a mechanistic framework in which the lunar gravitational signal is implemented via two independent layers — *behavioural-circadian* (phase-based) and *biomechanical* (wave-based, **ddF**) — converging at the level of **Ca^{2+}-dependent regulation** and the **HPA/HPG** neuroendocrine axes.

The obtained results substantiate the existence of a lunar tidal trigger as an independent population-level factor and open a new class of applied problems — from forecasting the load on emergency services to chronopharmacology — establishing a new class of applied tasks in chronobiology and population epidemiology.

Keywords: seleno-gravitational rhythm, SGR, lunar tidal jerk, ddF-detector, impulsive tidal wave, cross-city meta-analysis, HKSJ, PIEZO1, Ca^{2+} signaling, SCN, circadian regulation, neuroendocrine modulation, gravitational biology, behavior forecasting, gravitational chronopharmacology, mechanosensitive ion channels, population epidemiology, criminal records analysis, lunar gravity behaviour.

---

1. Introduction

1.1. Historical context and state of the literature

Although circalunar and circatidal rhythms are well documented in marine organisms and are considered an independent class of biological timers (Tessmar-Raible et al., 2011), ideas regarding the connection between the lunar cycle and the physiological and behavioral state of humans can be traced from Hippocrates to modern chronobiology. The principal results of the last two decades have moved this discussion from the realm of empirical observation into the realm of mechanistic biology: modulation of sleep architecture has been demonstrated under controlled conditions and in the absence of direct light exposure (Cajochen et al., 2013); synchronization of bipolar-disorder cycles with lunar gravitational rhythms has been documented across decades (Wehr, 2018); mechanosensitive ion

channels PIEZO1/PIEZO2 have been discovered (Nobel Prize, 2021; Coste et al., 2010), which respond to subtle changes in the mechanical state of tissues. Subsequently, their broad representation and functional significance in vascular, neuronal, and other mechanosensitive tissues have been established.

At the same time, the overwhelming majority of population-level studies of lunar influence are confined to a single methodological framework: comparing "lunar phase day *vs.* other days." The results of such an approach are, predictably, contradictory, and the cause of this is **methodological** rather than phenomenological.

Biological systems encode the **derivative** of a sensory stimulus rather than its absolute value: this is a basic principle of neurophysiology (the Weber–Fechner law, rapidly adapting receptors). The same principle holds for mechanoreception: PIEZO1 is a sensor of *changes* in membrane tension rather than of constant pressure (Li et al., 2014; Lewis and Grandl, 2015; Rode et al., 2017). Consequently, the biologically relevant parameter of lunar tidal forcing should not be the absolute force *F*, but rather its first and especially **second** derivative — the ***jerk* of the tidal potential.**

Prior to the present work, this approach had not been applied to large-scale social and medical data.

1.2. From personal observations to big data

The present study originated from a sequence of systematic personal observations of the author's own psychophysiological states and of the independent behavior of domestic and street animals over an extended period (more than a year, of which six months involved a regularly dated diary). The reproducible alignment of subjective and observable behavioral shifts with specific astronomical configurations gave rise to the working hypothesis: the lunar gravitational background contains **impulsive elements** registered by the nervous system as short-duration triggers, and these triggers should be visible in large-scale data reflecting the collective behavior and physiological state of a population.
A deliberate refusal of any retrospective "fitting" to lunar phases was a matter of principle, replaced by the development of a formal astronomical detector with an explicitly defined triggering criterion. The hypothesis thus formed was then put to the test against large open databases of U.S. police and medical services, which had no connection whatsoever to the author's personal observations.

1.3. Hypothesis and novelty

The hypothesis formulated in the present work. The lunar tidal potential contains at least two independent, biologically relevant layers: (i) a **static phase layer** — the traditional four lunar phases as a marker of long-period neuroendocrine modulation; and (ii) an **impulsive wave layer** — short-duration jumps in the second derivative d^2F/dt^2 acting as a biomechanical trigger for mechanosensitive tissues. These two layers have different

temporal profiles, different molecular targets, and should manifest themselves in different classes of biological and behavioral outcomes.

Elements of novelty of the present work:

1. **For the first time**, a detector of impulsive jumps in the second derivative of the tidal potential (hereafter — *ddF-detector* or *lunar tidal-wave detector*) has been developed and applied to social and medical time series at the city level.
2. **For the first time**, a cross-city meta-analysis of lunar influence on criminal and medical statistics has been conducted across five U.S. metropolitan areas using a random-effects model and HKSJ correction.
3. **For the first time**, the two layers of the lunar signal — phase-based and wave-based — have been formally separated and compared through an orthogonal decomposition of residuals (*Branch A vs. Branch B residualization*).
4. **For the first time**, the existence of a robust syzygy/quadrature dichotomy has been demonstrated on an array of more than 106 million records — a structural fingerprint with a clear neurophysiological interpretation.

1.4. Aims of the study

1. To develop and validate an original astronomical detector of impulsive lunar tidal waves based on the second derivative of the tidal potential, orthogonal to the static lunar phases.
2. To apply the detector to seven independent official databases of criminal and medical statistics from U.S. metropolitan areas, and to estimate effect sizes with conservative corrections for multiple testing and autocorrelation.
3. To carry out spatial (cross-city) and temporal (PRE / PRE2014) replication of the detected signals.
4. To describe the proposed neurophysiological and behavioral mechanism linking the gravitational trigger to the observed population-level shifts in the state of consciousness, behavior, and somatic health.

---

2. Data

2.1. Sources

Seven open official databases were used, accessed via the portals data.cityofnewyork.us, data.lacity.org, data.cityofchicago.org, opendataphilly.org, and data.sfgov.org. For each database, a unification procedure was carried out: aggregation by local calendar days, mapping of crime codes and medical dispatch codes, removal of duplicates and undocumented records. The resulting daily tables were brought to a unified format `date × crime_code → count`.

Table 1. Databases used.

Dataset	Period	Records	Link
Chicago Crimes	2001-01-01 — 2026-02-20	~7 000 000	Link
Los Angeles Crimes	2010-01-01 — 2022-12-31	~3 000 000	Link part 1 Link part 2
New York City Crimes	2008-01-01 — 2023-12-31	~7 300 000	Link
Philadelphia Crimes	2006-01-01 — 2026-02-26	~3 500 000	Link
San Francisco Crimes	2003-01-01 — 2017-12-31	~2 000 000	Link
New York City 911 Calls	2018-01-01 — 2025-12-31	~54 900 000	Link
New York City EMS	2005-01-01 — 2025-08-31	~28 700 000	Link
Total		**~106 400 000**	

2.2. Epochs

The analysis is stratified by the following epochs:

- **PRE2014** — data through 2014 inclusive. Used not as an independent cohort, but as a nested sensitivity sub-epoch, allowing the robustness of the gravitational signal to be assessed prior to the onset of the wave of decriminalization of offenses in several U.S. states.
- **PRE** — the full pre-pandemic period (through 1 March 2020). It serves as the principal analysis epoch and the most representative reference time series.
- **POST** — data after the pandemic (from 2022 onward). Excluded from the baseline pool due to shifts in societal behavior, but still processed by the algorithm and reported in CSV outputs.

The **COVID** epoch (2020-03-01 — 2021-12-31) is fully excluded from the analysis as a period of extreme regime change (lockdowns, altered mobility patterns, and changes in incident-reporting behavior).

2.3. Grouping of crime codes

Owing to the heterogeneous taxonomy of police codes across cities, two independent groupings were constructed:

- **Pure table** — individual codes of a single type of offense within each city (e.g., `PURE_BURGLARY`, `PURE_THEFT`, `PURE_ROBBERY`), enabling signal verification on maximally "clean" series.
- **Semantic families** — semantically grouped clusters (e.g., `AGGRESSION_TOTAL`, `PROPERTY_CRIME`, `SEXUAL_ASSAULT_CORE`), providing the required statistical power

through aggregation of semantically related codes. The full SEMANTIC_MAP is provided in the supplementary materials.

This dual grouping implements the principle of mutual verification: if a signal observed at the aggregated family level fails to reproduce in the pure samples, this would indicate a grouping artifact; the obtained results (§4.1) reveal no such artifact.

---

3. Methodology

3.0. Analytical architecture

The analysis is built upon two parallel and independent analytical pipelines. The **cross-city pipeline** (correlator → verdict) is applied to the five criminal databases and tests the same hypothesis across five independent cities, with aggregation via random-effects meta-analysis (HKSJ); the principal argument here is cross-city replication at minimal model complexity. The **single-database pipeline** (single_correlator → single_verdict) is applied to two unique databases for which external replication is not feasible (NYC 911, NYC EMS); here external replication is replaced by an internal robustness battery (block- and circular-permutation, bootstrap, slice stability, placebo, quasi-Poisson sanity-check). Both pipelines use the same astronomical engine (§3.1–§3.2), the same detrending procedure (§3.3), the same orthogonal Branch A / Branch B decomposition (§3.4), and the same TIER-label hierarchy (§3.10), differing only in the mode of inference and in the criteria for the final verdict.

3.1. Astronomical engine

All astronomical parameters were computed using the **Skyfield (v1.53)** library (Rhodes, 2019) and the **JPL DE421** ephemerides (Folkner et al., 2009), which provide sub-second accuracy relative to the timestamp. For each city, the reference instant was set to **local noon (12:00 local time)**, which ensures: (i) a unified daily reference time across all cities; (ii) a standardized day-to-day anchoring of the calculations within each city; (iii) correct alignment with the local calendar days into which the source data have been aggregated.

The full lunar–solar tidal scalar acting at the reference point is computed as

$$F_{\text{total}}(t) = \frac{M_{\text{moon}}}{r^3_{\text{moon}}(t)} + \frac{M_{\text{sun}}}{r^3_{\text{sun}}(t)} \cdot \mathcal{P}_2(\cos\theta(t)),$$

where M are the masses of the Moon and the Sun, r are the current geocentric distances, θ is the Moon–Earth–Sun angle, and

$$\mathcal{P}_2(\cos\theta) = \frac{3\cos^2\theta - 1}{2}$$

is the **second-degree Legendre polynomial** (the standard expansion of the tidal potential, providing the maximum at **syzygy** and the minimum at quadrature).

The parameters of the nodal (18.61-year) and apsidal (8.85-year) cycles are **computed and retained as separate contextual parameters**, `nodal_modulation ∈ [0.86, 1.14]` and `apsidal_modulation ∈ [0.97, 1.03]`.

The first and second derivatives of the tidal scalar are computed by finite differences with a step size of $\Delta t = 1$ h:

$$\frac{dF}{dt} \approx F(t) - F(t - 1\,\mathrm{h}), \qquad \frac{d^2F}{dt^2} \approx [F(t) - F(t - 1\,\mathrm{h})] - [F(t - 1\,\mathrm{h}) - F(t - 2\,\mathrm{h})].$$

For convenience of numerical representation and readability of thresholds, the values of dF/dt and d²F/dt² are hereafter reported in a form rescaled by multiplication by 10^6 (a dimensionless coefficient that does not affect statistical inference).

Reproducibility software environment. All computations were carried out in Python 3.10.11; key libraries: Skyfield 1.53, NumPy 2.2.4, pandas 2.2.3, SciPy 1.15.2, statsmodels 0.14.6. The full manifest is provided in `requirements.txt` of the repository.

3.2. ddF-detector of impulsive lunar tidal waves

Impulsive events — short-duration jumps in d^2F/dt^2 — are detected by a formal procedure with the following logic (the full implementation is in the repository, file `gen.py`):

"Down-start" event — an indicator of an approaching syzygy:

$$(x_{k-2} > -THR\_BASE) \ \wedge \ (x_{k-1} \le -THR\_BASE) \ \wedge \ (x_k \le x_{k-1} - THR\_STEP),$$

where `THR_BASE = 0.035`, `THR_STEP = 0.020` (values in units rescaled by multiplication by 10^6; thresholds were calibrated over a 24-year interval and fixed prior to the main run — see §6, artifact J). The condition corresponds to d^2F/dt^2 crossing the negative threshold with accelerated descent.

"Up-start" event — an indicator of an approaching quadrature (mirror logic):

$$(x_{k-2} < THR\_BASE) \ \wedge \ (x_{k-1} \ge THR\_BASE) \ \wedge \ (x_k \ge x_{k-1} + THR\_STEP).$$

Lookback / cooldown. In addition to the starting events themselves, "continuations" are flagged (`LOOKBACK_START = 4` days), conditional on a consecutive 3-point trajectory of dF/dt in the corresponding direction and the expiration of the cooldown period (`COOLDOWN = 2` days). In terms of the `gen.py` implementation, this means the presence of a starting event of the corresponding type within the last 4 **daily steps**, a monotonic 3-point trajectory of `dF_dt_micro` in the same direction, and the absence of an event of the same type within the last 2 **daily steps**. This provides a biologically reasonable registration profile — the event is flagged where tissues with inertia respond to the jump.

Each event is classified by the nearest lunar phase (**syzygy** or **quadrature**) and by relative position (**before** or **after**), yielding seven binary triggers:

- `wave_event` ; `wave_syzygy` ; `wave_quadrature` ;

- `wave_before_syzygy` ; `wave_before_quadrature`
- `wave_after_syzygy` ; `wave_after_quadrature`

A key empirical property of the detector: in the vast majority of cases the event is registered within a window of **±1…2 days** of the lunar phase and very rarely coincides with the phase itself. This implies that wave triggers are constructively orthogonal to static phase labels and can be used as an independent signal.

Biological rationale for selecting *d²F/dt²* as the working variable:

- **Neurophysiologically:** many sensory systems primarily encode the change of a stimulus; PIEZO channels, in particular, respond to changes in mechanical tension rather than to its prolonged static maintenance (Li et al., 2014).
- **Mathematically:** extrema of $\frac{dF}{dt}$ correspond to zeros of $\frac{d^2F}{dt^2}$, whereas peaks of $\left|\frac{d^2F}{dt^2}\right|$ pick out the moments of the sharpest change in the rate of variation of the external field itself.
- **Geophysically:** the magnetohydrodynamic signal of coastal tidal currents is related to the velocity of conducting seawater motion in the Earth's magnetic field, thereby emphasizing the physical relevance not only of the static tide level but also of its dynamics (Sabaka et al., 2016).

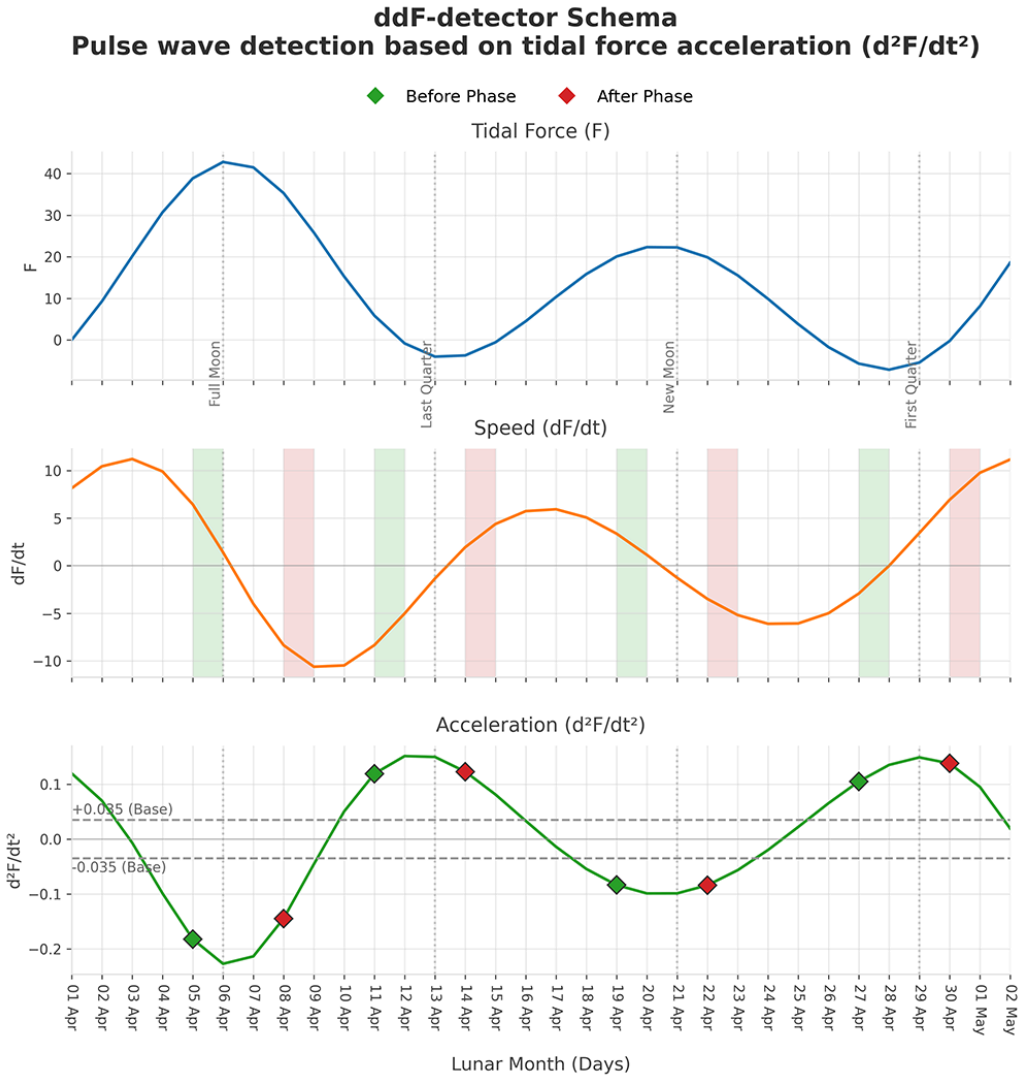

3.3. Detrending and deseasoning

Prior to effect estimation, each time series of `log1p(counts)` for the pair `(city × epoch × crime_type)` is processed by OLS regression with the following set of predictors:

- a polynomial trend of degree 4 in time;
- an annual Fourier basis (3 harmonics);
- day-of-week dummies (6);
- U.S. federal holidays plus ± 1-day lags around each holiday.

The resulting residuals are then subjected to symmetric winsorization (`1% / 1%`) for robustness against extreme outliers (mass events, major incidents), suppressing the influence of tail values without introducing systematic bias.

3.4. Orthogonal Branch A / Branch B decomposition

A principal methodological step of the present work is the separation of the analysis into two orthogonal branches:

Branch A *(used for `PHASE_` triggers).**
OLS from §3.3 without the inclusion of any lunar variables. The residuals retain all lunar-associated variance not explained by non-lunar covariates and allow a clean estimation of the contrast **"lunar phase days vs. days outside the ±1-day phase window."**

Branch B *(used for `WAVE_* triggers).* The OLS includes the full set of static lunar dummies – Full / New / First / Last` with lags $\{-1, 0, +1\}$ (12 variables). The residuals have the static phase component partialled out, and the estimation of wave effects in these residuals answers a clear question: **"Does the impulsive ddF wave add information beyond the static lunar phase?"**

Since wave events are by construction often located within a window of $\pm 1 \ldots 2$ days of the lunar phase, and Branch B explicitly includes static phase dummies for lags $\{-1, 0, +1\}$, these regressors inevitably absorb part of the true wave variance, particularly for events close to the phase itself. Consequently, all estimates of `WAVE_*` effects in the present work are **lower bounds (*lower bounds*) of the true wave signal.** This conservatism is a deliberate methodological choice: it is better to underestimate the effect than to attribute to the wave what is already explained by the known static phase.

3.5. Effect size and its variance

As the primary effect-size measure we employ **Hedges' g** (Hedges, 1981) — a standardized difference of mean residuals between "exposed" (target) and "control" days with the Hedges (J) small-sample correction:

$$g = J \cdot \frac{\bar{r}_{\text{target}} - \bar{r}_{\text{control}}}{s_{\text{pooled}}}, \qquad J = 1 - \frac{3}{4(n_{\text{target}} + n_{\text{control}}) - 9},$$

where $\bar{r}$ are the means of the winsorized residuals (§3.3) over the respective subsets of days, and s_{pooled} is the pooled standard deviation. The variance of the effect size Var(*g*), required for the subsequent random-effects meta-analysis, is also computed analytically:

$$\mathrm{Var}(g) = \frac{n_{\text{target}} + n_{\text{control}}}{n_{\text{target}} \cdot n_{\text{control}}} + \frac{g^2}{2(n_{\text{target}} + n_{\text{control}})}.$$

As a nonparametric backstop test, a two-sided Mann–Whitney U-test is additionally computed. This is especially important for the single databases (NYC 911, NYC EMS), where cross-city meta-aggregation is not available: in the single-database pipeline the nonparametric criterion `p_MW` serves as the baseline significance condition, after which a signal is admitted to the TIER classification (§3.10) only upon additional confirmation of internal robustness — through permutation tests, slice stability, and placebo control.

In addition, for each significant result, a relative shift is computed on the raw counts:

$$\Delta\% = 100 \cdot \left(e^{\bar{r}_{\text{target}} - \bar{r}_{\text{control}}} - 1\right),$$

allowing the standardized effect to be translated into physically interpretable percentages of excess over background.

3.6. Lag-window scheme

The biological response to any trigger need not appear strictly on the day of the trigger. Vascular, neuroendocrine, and behavioral reactions to a gravitational impulse may shift by a day in either direction — owing to the inertia of the HPA axis, nighttime sleep, and behavioral delay between an internal state and a registered event.
To account properly for this jitter, each effect was estimated independently within four time windows:

- `lag = -1` — one day before the event;
- `lag = 0` — on the day of the event;
- `lag = +1` — the day after;
- `lag = WIN3` — a 3-day exposure window, the union of `{-1, 0, +1}`.

WIN3 is the principal timing-robust estimate of the effect. Unlike the per-event argmax approach (selection of the "best" lag for each event), a fixed three-day window eliminates selection bias and the so-called *winner's curse* — artifactual amplification of the estimate due to optimization over a family of windows. The individual lags {−1, 0, +1} are reported as a diagnostic lag profile supporting biological interpretation (e.g., "the effect concentrates exactly on the day of the event" versus "the effect is spread over two days").

The window is deliberately limited to ±1 day: a wider span (±2…3 days), given the synodic period of ~29.5 days, would begin to overlap with neighboring astronomical triggers (e.g., the intervals phase ↔ wave), which would destroy orthogonality.

3.7. Robustness battery (single-database pipeline)

For the two unique single databases — **NYC 911 dispatches** and **NYC EMS** — it is impossible to construct a cross-database meta-aggregation: comparable public services with identical codes in other cities are not available in accessible form. This data property fundamentally affects the inferential strategy: the absence of external replication is compensated by an **internal robustness battery**, comprising several independent procedures, each targeting its own class of potential artifacts (full implementation — `single_correlator.py`).

By default, the extended robustness battery is computed for `lag = 0` and `lag = WIN3`; for lags `-1` and `+1`, it can be enabled separately via configuration.

(i) **Block-permutation tests** with block lengths of 14 and 28 days.
In the implementation, blocks of the residual series are permuted, while the masks of exposed and control days are held fixed. This preserves the local autocorrelation structure within a block while destroying the original temporal alignment of the effect with the trigger. The empirical `p`-value is defined as the proportion of permutations for which $|g_{\text{null}}| \geq |g_{\text{obs}}|$.

(ii) **Circular-shift permutation** with a minimum shift of ≥ 7 days.
The residual series is cyclically shifted relative to a fixed exposure mask, preserving the internal structure of the series but destroying their mutual phase alignment. This is a conservative test against the hypothesis of random coincidence of seasonal or autocorrelated processes.

(iii) **Bootstrap 95% CI** for Hedges' g (percentiles 2.5/97.5).
Used as an additional measure of effect stability with respect to sampling variability.

(iv) **Slice stability** across 5 contiguous time slices.
Within each slice, the effect is estimated separately, after which `slice_sign_frac` — the fraction of valid slices in which the sign of the effect matches the overall direction of the signal — is computed. This test rules out findings concentrated in only one local epoch.

(v) **Placebo resampling**.
Random subsamples of size `n_target` are repeatedly drawn from the set of control days; the proportion of such "pseudo-targets" yielding $|g| \geq |g_{\text{obs}}|$ defines the empirical `placebo-p`. This procedure addresses the class of artifacts in which "given sufficiently large `n`, a significant effect can be obtained on almost any random subsample."

(vi) **Quasi-Poisson GLM** on raw counts as a *sanity check*.
For effects passing the main estimation stage, the sign and significance of the trigger coefficient are additionally checked in a quasi-Poisson model on the raw counts. The output tables also record an indicator `model_disagree` if the sign of this coefficient disagrees with the sign of Hedges' g.

In the single-database verdict, the final **TIER classification** is constructed from `p_MW`, `slice_sign_frac`, the number of successfully passed permutation tests, and `p_placebo_emp`; the bootstrap CI and quasi-Poisson GLM serve as additional diagnostic checks. The full thresholds for the TIER classification are described in §3.10.

3.8. Multiple-comparison control and gatekeeping

Multiple-testing control is performed via the **Benjamini–Hochberg (FDR)** procedure (Benjamini & Hochberg, 1995) within strata `(epoch × trigger_family × lag)` across all event types within a single analytical unit — a city in the main pipeline and a database in the single-database pipeline. Such stratification is essential: it avoids pooling into one common correction the statistically distinct families of tests pertaining to different trigger hypotheses and time windows.

For each tested combination `(city/database × epoch × crime_type × trigger × lag)`, a three-tier **gatekeeping** rule applies:

- mean number of events of the given type `≥ 8/day`;
- `n_target ≥ 15` exposed days;
- `n_ctrl ≥ 60` control days.

These thresholds are set identically in both computational branches and serve as the minimum admission condition for a row to enter the main effect-size estimation. If at least one of them is violated, the combination is not analyzed as a full-fledged signal and is assigned the technical status `INSUFFICIENT_N`.

3.9. Cross-city meta-analysis (random-effects, HKSJ)

Cross-city aggregation is performed in `verdict.py` within a **random-effects meta-analysis** over the effect-size rows supplied by `correlator.py`. The meta-analysis is built on two levels: **family-level** (semantically grouped crime families) and **crime-type level** (a diagnostic table by individual crime types). At the family level, a single city may contribute several semantically mapped `crime_type` effects, so the output tables separately record `k_total` (the total number of effect-size rows) and `k_cities` (the number of unique cities).

The true effect is treated as a random realization from a distribution with a global mean and a between-study variance `τ²`. The parameter `τ²` is estimated iteratively in the scheme denoted in the code as **Empirical-Bayes / REML τ² estimation**, after which the final meta-estimate is computed with weights `1 / (Var(g) + τ²)`.

A key methodological choice is the **HKSJ correction** (Hartung–Knapp–Sidik–Jonkman). In `verdict.py`, it is implemented through a correction of the standard error of the meta-estimate and the use of the **t-distribution** with `df = k − 1` for the computation of `p_meta` and the 95% confidence interval. This choice is especially important when the number of cohorts is small and renders the inference deliberately more conservative, reducing the risk of an anti-conservative underestimation of the standard error in the presence of non-zero

between-city heterogeneity (Hartung & Knapp, 2001; Sidik & Jonkman, 2002; IntHout et al., 2014).

Heterogeneity is assessed via Cochran's `Q` statistic, its `p`-value (`p_Q`), and the `I²` share (Higgins & Thompson, 2002). When `I² = 0%`, the observed between-row variability does not exceed the expected within-city sampling variability; as `I²` increases, the contribution of between-city heterogeneity of the effect grows.

3.10. TIER classification

For the final summary, signals are classified on a strict scale.

Cross-city (family-level meta-analysis):

- `TIER1_PUBLISH` — `p_meta < 0.005` ∧ `k_cities ≥ 3` ∧ `I² < 40%`.
- `TIER2_PROMISING` — `p_meta < 0.010` ∧ `k_cities ≥ 3` ∧ `I² < 60%`.
- `TIER3_SIGNAL` — `p_meta < 0.050` ∧ `k_cities ≥ 2`.
- `NULL` — otherwise.
- `N/A` — secondary trigger or non-target epoch.

Single-database (internal robustness):

- `TIER1_PUBLISH` — `p_MW < 0.005` ∧ `slice_sign_frac ≥ 0.80` ∧ `3/3 permutation pass` ∧ `placebo pass`.
- `TIER2_PROMISING` — `p_MW < 0.010` ∧ `slice_sign_frac ≥ 0.75` ∧ `≥ 2/3 permutation pass` ∧ `placebo pass`.
- `TIER3_SIGNAL` — `p_MW < 0.050` ∧ `slice_sign_frac ≥ 0.50`.
- `DISCARDED_UNSTABLE` — insufficient slice stability; specifically, a row is discarded immediately when `slice_sign_frac < 0.40`, and is also not admitted even into `TIER3_SIGNAL` if, with `p_MW < 0.050`, the condition `slice_sign_frac < 0.50` holds.
- `NULL` — otherwise.
- `N/A` — secondary trigger or non-target epoch.

In `single_verdict.py`, the semantic mapping is used to decode `signal_title` and to group known signals, but rows that do not fall into the mapping are retained and analyzed as `RAW_<crime_type>`. Consequently, the single-database verdict is not limited only to pre-defined semantic groups.

As the headline lag **for the principal TIER1 result** of the present work, WIN3 is used as the timing-robust estimate. For additional diagnostic clusters (§4.4), the individual lags {−1, 0, +1} are additionally reported, highlighting the specific temporal profile of the effect (e.g., concentration on the day of the event itself vs. a pre-/post-impulse shift).

4. Results

4.1. Overall pipeline statistics

The final analysis was supplied with a total of **10,192 valid rows** from the meta-pipeline (5 cities, 13 semantic families + pure tables) and **19,916 rows** from the single-database pipeline (NYC 911 + NYC EMS, across all codes that passed the gatekeeping criteria of §3.8). Of the 1144 family-level rows for primary triggers in the PRE and PRE2014 epochs, **1,127** were classified as NULL; of the remaining 17 rows: 1 TIER1, 1 TIER2, 15 TIER3 (§4.2–§4.4) — that is, the overwhelming majority of combinations (`family × trigger × epoch × lag`) does not show a significant signal, which is itself an important observation: the method **does not "find" effects everywhere**, as would be the case under artifactual sensitivity.

4.2. Principal finding: PROPERTY_CRIME × WAVE_before_quadrature (TIER1)

The only result passing the TIER1 criterion in the cross-city meta-analysis is the reduction of property crime (`PROPERTY_CRIME`) within the three-day impulsive-wave window *before* the lunar quadrature:

> **PROPERTY_CRIME × WAVE_before_quadrature × PRE2014 × WIN3:**
> g_meta = **−0.032**, 95% CI [−0.052, −0.011], p = **3.93×10^{-3}**, I^2 = **0%**, k = 21 effects across 5 cities.

Here k = 21 is the total number of effect-size contributions that entered this meta-row from 5 cities: Chicago = 3, LA = 7, NYC = 4, Philadelphia = 4, SF = 3.

Contributions from individual cities: Chicago [−0.026, −0.028, −0.035]; Los Angeles [seven estimates, median ≈ −0.04]; New York [−0.015 … −0.077]; Philadelphia [−0.034 … −0.094]; San Francisco [+0.020, −0.023, −0.067].

Within individual cities there are also locally divergent contributions, but at the level of cross-city family aggregation the final effect remains negative and statistically significant. **The value I^2 = 0%** in this TIER1 result means that the entire variance of the observed per-city effects is explained by within-city sampling variability, and no detectable between-city heterogeneity of the effect remains. Critically, this result is obtained **without post hoc** exclusion of divergent local contributions from the family composition; consequently, the signal is not reducible to manual fitting to a single sign.

The same signal is reproduced in the full PRE epoch:

> **PROPERTY_CRIME × WAVE_before_quadrature × PRE × WIN3:**
> g_meta = −0.023, 95% CI [−0.039, −0.007], p = 7.4×10^{-3}, I^2 = 0%, k = 21 — **TIER2_PROMISING (temporal replication)**.

This means that the signal detected in the 2001–2014 sub-cohort is reproduced when the temporal window is extended to the full pre-COVID epoch (2001 — March 2020). Formally, PRE2014 is a nested sub-cohort within PRE; however, the additional years (2015 — March 2020), included only in PRE, span an independent temporal segment of American criminal history, distinct in political, demographic, technological, and reporting structure — which makes the coincidence in sign and magnitude of the effect a reliable equivalent of partial temporal replication.

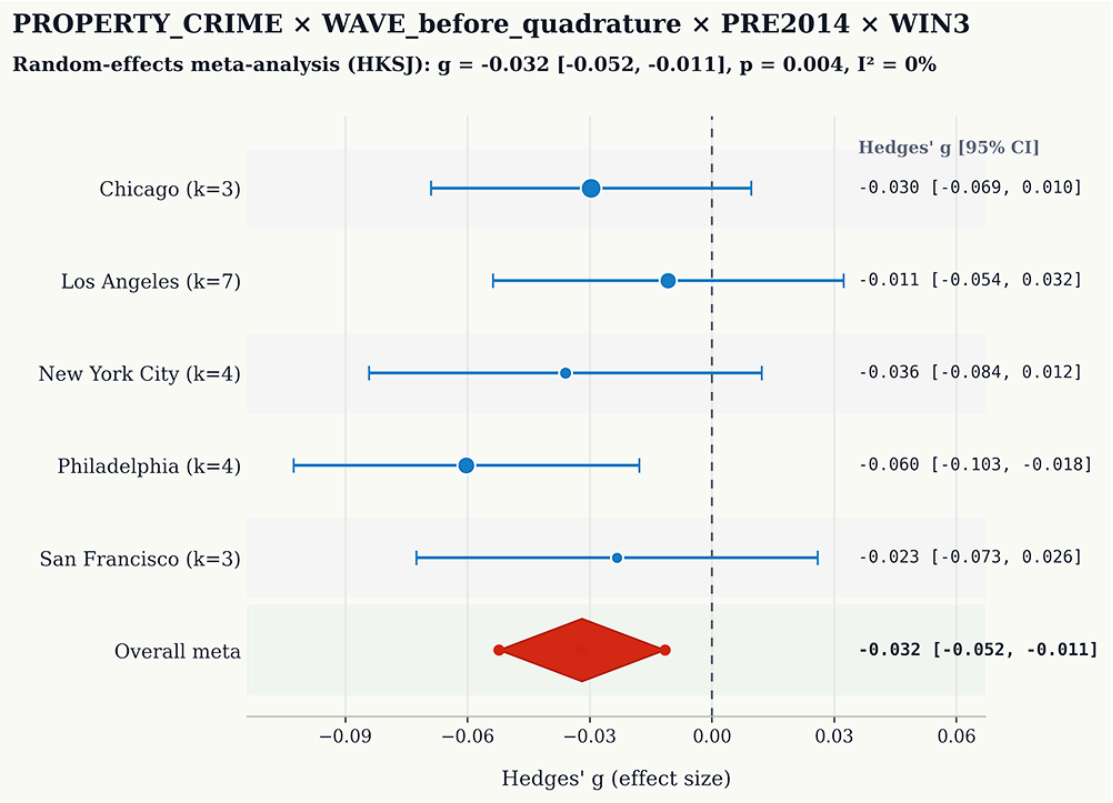

4.3. Effect size — interpretation on the correct scale

The magnitude g_meta ≈ −0.03 is objectively small on the absolute scale of Cohen's d (by convention, d around 0.2 is considered small). However, this estimate must be interpreted in the context of the **noise environment** of metropolitan areas and the **conservative design of the analysis itself**:

1. Daily counts of property crimes in U.S. metropolises vary by tens of percent between ordinary weekdays owing to weather, events, patrol policy, and reporting. Extracting a stable shift on the order of a few percent, synchronized with an astronomical trigger, from such noise is a non-trivial task.
2. WAVE_ signals are estimated on the Branch B residuals, where static phase dummies have already absorbed part of the true wave variance (§3.4). Consequently, the estimate of −0.032 should be interpreted rather as a **lower bound of the observed effect** than as its exhaustive approximation.

3. The HKSJ correction at $k = 5$ cities imposes an additional penalty on the confidence interval. A signal that survives HKSJ at such k, is reproduced across 5 cities and in two temporal epochs with $I^2 = 0\%$, cannot be interpreted otherwise than as a **real, small-amplitude, population-level shift**.

In other words, the principal argument in favor of the reality of the effect is not its size, but the **combination of robustness + cross-city homogeneity + temporal replication** in a noisy urban environment and a deliberately conservative model. This is the "signal fingerprint" in the sense of signal theory.

4.4. Additional cross-city results (TIER3) and their structure

Beyond the TIER1/2 signal described above, the cross-city meta-analysis records **15 TIER3 results**. Their distribution by families and triggers is not random and forms three clear clusters:

"Property bipolar" cluster. PROPERTY*CRIME, PURE_VEHICLE_THEFT, PURE_ROBBERY systematically go **negative** on WAVEquadrature and PHASE-*Quadrature/FirstQ/LastQ — and **positive** on WAVE_syzygy / WAVE_before_syzygy:

- PROPERTY_CRIME × WAVE_syzygy × PRE × lag=0: $g = +0.021$, $p = 3.2 \times 10^{-2}$;
- PROPERTY_CRIME × WAVE_before_syzygy × PRE × lag=0: $g = +0.028$, $p = 3.4 \times 10^{-2}$;
- PROPERTY_CRIME × PHASE_LastQ × PRE2014 × lag=0: $g = -0.055$;
- PURE_ROBBERY × WAVE_before_quadrature × PRE2014 × lag=0: $g = -0.085$;
- PURE_VEHICLE_THEFT × PHASE_FirstQ × PRE × lag=−1: $g = -0.100$.

Property crime behaves as a **bipolar oscillator** synchronized with the lunar cyclic potential: an increase in syzygy windows and a decrease in quadrature windows.

"Aggression on New Moon" cluster. AGGRESSION_TOTAL produces a consistent positive sign on the static new moon and on the overall syzygy in both epochs:

- AGGRESSION_TOTAL × PHASE_New × PRE × lag=0: $g = +0.072$, $p = 3.4 \times 10^{-2}$;
- AGGRESSION_TOTAL × PHASE_New × PRE2014 × lag=0: $g = +0.089$, $p = 3.6 \times 10^{-2}$;
- AGGRESSION_TOTAL × PHASE_Syzygy × PRE2014 × lag=0: $g = +0.060$, $p = 4.8 \times 10^{-2}$.

This no longer looks like an isolated random TIER3 episode: the same "family × trigger" pair recurs across two epochs with the same direction of effect. Therefore, this result should be interpreted as a consistent recurring pattern, rather than a single chance crossing of the $p < 0.05$ threshold.

"Pure verifies aggregated" cluster. PURE samples (single codes) confirm the sign of the aggregated family signals:

- PURE_ROBBERY and PURE_VEHICLE_THEFT reproduce the negative sign of PROPERTY_CRIME in the quadrature window;
- PURE_VANDALISM, PURE_BURGLARY, and PURE_THEFT yield consistent directions at weakened significance (a characteristic feature of strict pure samples with limited statistical power).

This substantially weakens the standard objection that the effect is allegedly generated entirely by artificial aggregation of dependent variables into families: upon switching to narrower PURE families, the direction of the signal is preserved.

4.5. Results on the single databases (NYC 911 and NYC EMS)

In the single-database pipeline, **14 TIER1 + 17 TIER2 + 115 TIER3** signals were recorded. The complete tables are provided in the supplementary materials; what follows is a structural summary along the bio-behavioral axes.

(A) Psyche and behavior (AGGRESSION, EDP, OTHER EMS).

- **OTHER EMS (Unclassified EMS call)** × WAVE_before_quadrature × WIN3: $g = +0.103$, $p = 2.7 \times 10^{-3}$ in the PRE epoch and $g = +0.121$, $p = 4.0 \times 10^{-3}$ in the PRE2014 epoch — **TIER1 in both epochs**. This is a wave-driven, biomechanical, non-specific field — an increase in the load on the ambulance service that **is not reducible to a specific diagnosis** — the principal "wave-driven" non-specific signal in the entire dataset.
- **Request for assistance to a uniformed patrol officer on the street** × WAVE_after_syzygy × lag = 0: $g = +0.371$, $p = 1.8 \times 10^{-3}$ — TIER1. A marker of street "blowups" after the full moon / new moon.
- **Suspicious object in the subway / transit** × WAVE_event × lag = 0: $g = +0.245$, $p = 2.8 \times 10^{-3}$ — TIER1. And the same signal × WAVE_after_syzygy: $g = +0.331$ — TIER2. A marker of collective anxiety.
- **EDP (Emotionally Disturbed Person) indoors** × WAVE_before_quadrature: $g = -0.214$, $p = 5.4 \times 10^{-3}$ — TIER2. A paradoxical **decrease** in indoor psychiatric calls before quadrature. On the same trigger, but in the "public transport" scene — $g = +0.363$ (TIER3). The scene determines the direction.
- **Bladed-weapon conflicts** in syzygy windows: reproducible positive signals on WAVE_after_syzygy and WAVE_syzygy.

(B) Cardiovascular system.

- **HYPTN (hypertensive crisis)** × PHASE_LastQ × WIN3: $g = +0.156$, $p = 1.8 \times 10^{-3}$ — TIER1.
- **Cardiac indoors (911)** × PHASE_FirstQ × WIN3: $g = +0.417$, $p = 1.1 \times 10^{-3}$ — TIER1.
- **CARD (Cardiac Arrest)** × PHASE_Full × lag = 0: $g = -0.193$, $p = 8.0 \times 10^{-3}$ — TIER2. Reproduced in PRE2014 on PHASE_Syzygy: $g = -0.166$ — TIER3. **A decrease in cardiac arrests on the full moon / syzygy** — a stable effect recurring across epochs.

- **CVA (Stroke)** × WAVE*event* × *PRE2014: g = +0.082, p = 4.0 × 10⁻² — TIER3. The signal is confined to the wave branch (Branch B): the corresponding PHASE* triggers for CVA in the same epoch do not reach the TIER3 threshold, indicating specificity of the response to the impulsive (rather than static) component of the lunar potential.
- **INBLED (internal GI bleeding)** × WAVE_after_quadrature × lag = 0: g = +0.146 in PRE2014 (TIER2) and g = +0.111 in PRE (TIER3) — **replication of sign and magnitude across two epochs**.

Thus, for the heart and vessels, a robust **syzygy / quadrature dichotomy** emerges: lunar syzygies (especially the full moon) are accompanied by a *decrease* in acute cardiac events, whereas quadratures and post-wave windows are accompanied by an *increase* in acute vascular events (hypertensive crisis, stroke, GI bleeding).

(C) Obstetrics.

- **OBLAB (Labor)** × PHASE_FirstQ × lag = 0: g = +0.201, $p = 3.9 \times 10^{-3}$ — TIER1. A static phase signal of labor activity.
- **OBMIS (Threatened miscarriage)** × WAVE_before_quadrature × lag = 0: g = −0.172 in PRE2014 (TIER2), g = −0.117 in PRE (TIER3) — **replication across two epochs**. The signal is specifically wave-based: the preceding quarter phase does not reproduce the same effect.

(D) Respiration and autonomic disturbances.

- **RESPIR (Difficulty breathing)** × PHASE_New: g = +0.227 (PRE2014, TIER3) and g = +0.128 (PRE) — replication of sign.
- **UNC (Loss of consciousness)** × PHASE_New × PRE2014 × lag = 0: g = +0.197 — TIER2.

(E) Road traffic injuries.

- **PEDSTR (pedestrian struck by a vehicle)** yields consistent signs in the syzygy / WAVE_event windows with a marked increase in the rising phase of the impulse and a decrease directly on the new moon (a likely effect of reduced pedestrian flow during the dark night).

4.6. Synthesis: state windows

Combining all TIER1/TIER2/TIER3 signals yields a coherent picture in which each window of the lunar cycle has a characteristic bio-behavioral "fingerprint". This picture is a product of the data, not of theory; it emerged post hoc when the results were sorted by trigger type.

Table 2. Characteristic dominant signals across the windows of the lunar cycle.

Window	Behavioral layer	Biomechanical / somatic layer
before_syzygy (−1…−2 days)	Street aggression ↑, weapons on the street ↑, RTAs ↑, property ↑	MVAINJ ↑ (RTA injuries), pedestrian injuries ↑
PEAK New Moon	Aggression (meta) ↑, indoor violence ↓	RESPIR ↑, UNC ↑, syncopes ↑
after_syzygy (+1…+2 days)	Assistance calls to patrol ↑, knife conflicts ↑	Cardiac on the street ↑, residential fires ↑
PEAK Full Moon	Residential alarms ↑, violence ↓, weapons ↓	CARD/ARREST ↓, SHOT ↓
before_quadrature (−1…−2 days)	Property ↓, indoor EDP ↓, conflicts ↓; EDP in transit ↑	OBMIS ↓ (Threatened miscarriage), OTHER EMS ↑ (non-specific load), STAB ↓ (Stabbing)
PEAK First Quarter	Apartment burglaries ↑, indoor activity ↑, weapons on the street ↓	OBLAB ↑ (labor), Cardiac indoors ↑
PEAK Last Quarter	Indoor evacuations ↑, premises entries ↑, EDP ↓, indoor aggression ↓	HYPTN ↑ (blood pressure), vascular load ↑
after_quadrature (+1…+2 days)	—	INBLED ↑ (internal GI bleeding)

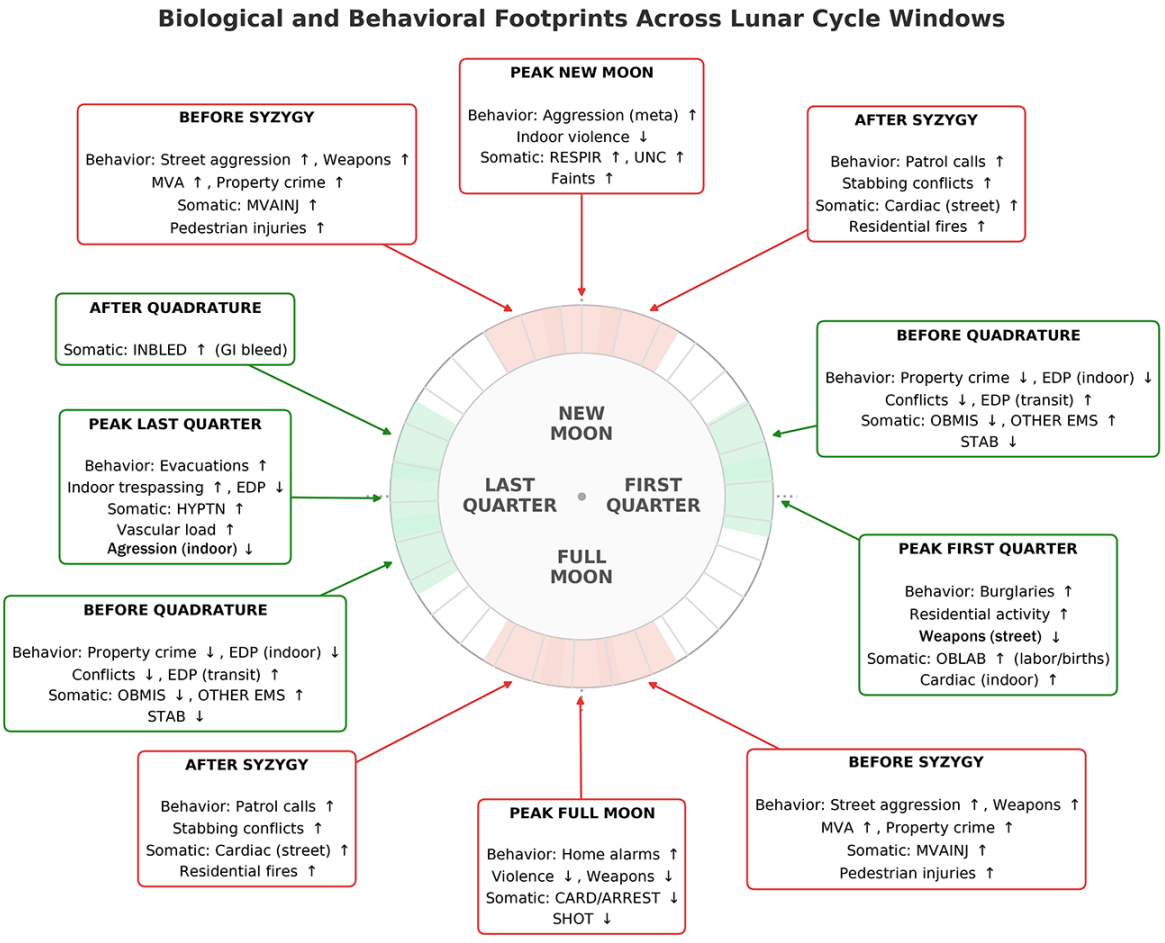

Of particular note is the **syzygy/quadrature dichotomy** as a cross-cutting structural fingerprint of the data:

- for the heart: syzygy ↓ acute arrests, quadrature ↑ acute vascular events;
- for property: syzygy ↑, quadrature ↓;
- for psychiatry: phases — fewer call-outs, quadrature — especially distinctly downward; the sole exception is public transit.

Such a consistent dichotomy cannot be explained by random noise: it requires a **physical mechanism acting differently in syzygy and quadrature windows**. The corresponding mechanistic framework is considered in §5.1.

4.7. POST epoch and extended tables

POST-epoch data (from 2022, the post-COVID period) are processed by exactly the same analytical pipeline at the level of code, yet are not used to support conclusions. Distortion of the signals in POST is to be expected and is mathematically explainable: (i) the sharp reduction of temporal coverage and of the size of the control sample (absence of several cities), and (ii) the lingering aftereffects of the 2020–2021 regime shifts, which broke the classical patterns of street activity due to remote work.

Since the use of the POST period entails unacceptable risks from social confounders, this epoch cannot serve as a valid sensitivity stratum. Nevertheless, the verdict scripts produce a full-fledged meta-analysis for the POST epoch: the results are automatically saved without the assignment of publication-grade significance levels (TIER) and are available to researchers in the files `outputs/verdicts/sgr_meta_family.csv` and `outputs/verdicts/sgr_single_verdict_signals.csv`. The complete family groupings of codes (SEMANTIC_MAP, 13 families and pure tables) are provided in Appendix A of the repository (`docs/semantic_map.md`).

---

5. Discussion

5.1. Two independent layers of the lunar signal

The central substantive observation of the present study: lunar tidal influence on population-level indicators is registered in the data not as one but as **two independent signal axes**:

(I) Behavioural-circadian (phase) layer. Acts on the time scale of a day to a week, anchored to the static calendar phase of the Moon, and manifests in indicators related to behavioral control, sleep, and neuroendocrine background: aggression on the new moon, anxiety episodes on the full moon, obstetric activity on the first quarter. Molecular targets — the melatonin axis (via ipRGCs in humans exposed to outdoor nocturnal light), the HPA and HPG axes, and the circadian pacemaker of the suprachiasmatic nucleus.

(II) Biomechanical (wave, ddF) layer. Acts via short-duration impulses of the second derivative of the tidal potential, manifests in indicators related to mechanosensitive tissues and rapid vascular events: stroke, internal GI bleeding, threatened miscarriage, non-specific load on the ambulance service. The molecular target — the PIEZO1/PIEZO2 family of mechanosensitive channels and downstream Ca^{2+} signaling.

Empirical evidence of the orthogonality of the two layers is provided by the Branch A / Branch B methodology (§3.4). The WAVE *signals are estimated on residuals* ***from which the static lunar phases have already been removed*** *as explanatory variables. The fact that WAVE* effects persist (and reach TIER1 in the cross-city meta-analysis) means that the ddF wave impulse adds information **beyond** the static phase. Were the two signals the same effect re-encoded in two ways, Branch B would have nullified WAVE_*.

This does not happen. Moreover, the signals of the two layers distribute across **different classes of outcomes**:

- phase signals concentrate in behavioral and circadian-dependent indicators (aggression, anxiety episodes, the cardiac quiescence of the full moon, labor activity);
- wave signals — in acute biomechanical and vascular events (stroke, internal bleeding, non-specific ambulance load, threatened miscarriage).

This distribution is biologically coherent: it reflects different temporal scales of response and different molecular pathways.

5.2. Physical channels of lunar influence

Between the astronomical lunar–tidal potential and the cellular response stands a chain of physical channels, each with its own temporal profile and molecular target. Six of them have to date undergone independent verification in the peer-reviewed literature and form the working mechanistic framework:

(A) Direct gravitational channel. The change in the magnitude of gravitational acceleration Δg between apogee and perigee reaches ~12% of the lunar component; the corresponding rhythmic modulation of the vascular mechanical environment is translated into an endothelial response via the mechanosensitive PIEZO1 channel, which senses shear stress and triggers a Ca^{2+}-dependent signaling cascade involving ATP, eNOS, and vasoactive pathways (Li et al., 2014; Wang et al., 2016; Rode et al., 2017). Additional support for this interpretation comes from the fact that, under simulated microgravity, endothelial cells exhibit increased PIEZO1 expression and a PIEZO1-dependent Ca^{2+}-mediated response, confirming the sensitivity of this channel to gravitational–mechanical environmental shifts (Wang et al., 2024). This is the principal channel of the biomechanical (wave) layer.

(B) Magnetospheric–electric channel. The magnetosphere, the ionosphere, and the near-surface atmospheric layer are linked through the global atmospheric electrical circuit, within which variations of the space and geophysical environment are capable of modulating electrical conditions near the Earth's surface (Rycroft et al., 2008). On the full moon, the

Moon passes through the Earth's geomagnetic tail (magnetotail), where electromagnetic disturbances are recorded that are considered a possible lunar-dependent external modulator of this channel; in particular, elevations of electric fields above the usual background and a potential influence on calcium signaling and other electrosensitive biological processes are discussed (Bevington, 2015).

(C) Light channel. Under conditions of natural exposure, ipRGCs of the retina are capable of detecting weak nighttime light, including the light of the full moon, and of suppressing nocturnal melatonin via the retinohypothalamic tract (Lucas et al., 2014); the typical illuminance of the full moon is approximately 0.05–0.1 lux, reaching about 0.3 lux only under near-ideal conditions (Kyba et al., 2017). In a metropolis, where the nocturnal street illuminance is typically orders of magnitude higher than lunar, the contribution of this channel is sharply attenuated, which makes it unlikely that the effects observed in the present work can be explained by lunar light alone.

(D) Ionospheric channel. The lunar tidal signal in the ionosphere (the M2 component) is modulated by the anomalistic month; the maximum of its amplitude in GNSS TEC is observed with a lag on the order of 3 days after perigee, which is consistent with the time of vertical propagation of the tidal wave from the lower atmosphere to the dynamo region of the ionosphere (Hocke, 2025). Geomagnetic stimuli of Earth-strength can elicit a reproducible alpha-band response in the human brain (Wang et al., 2019). This channel is a candidate for explaining the delayed and cumulative components of the observed signals.

(E) Magnetohydrodynamic channel of coastal tides. The motion of electrically conducting seawater in the geomagnetic field generates magnetohydrodynamic currents and secondary magnetic fields with amplitudes on the order of a few nT, registered in satellite data (Sabaka et al., 2016; Grayver & Olsen, 2019). Since this signal is physically defined by the motion of water, it constitutes an independent geophysical argument that not only absolute levels of tidal forcing but also its dynamical derivatives are biologically and geophysically meaningful.

(F) Infrasonic channel. Syzygy configurations are accompanied by enhanced tidal deformation of the Earth–ocean–atmosphere system and may increase the likelihood of weak microseismic and infrasonic disturbances. The neural effect of near-threshold infrasound on the anterior cingulate cortex and the right amygdala has been documented in a double-blind fMRI study (Weichenberger et al., 2017), which makes such a channel biologically plausible.

The key point: all six channels converge on **Ca^{2+}-dependent regulation**, although the degree of this linkage varies — from direct mechanosensitive input to neuroendocrine and circadian mediation. This naturally explains why the signal in the data turns out to be distributed across several independent physiological axes rather than localized in one, and why the statistical effects are small in amplitude yet stable in sign and cross-city replication: they are the sum of small contributions from several parallel pathways with a common final mediator.

The specific contribution of each channel to the observed effects is an open question; the present work does not claim to decompose them, but rather fixes the integrated signal.

5.3. The biomechanical axis: PIEZO → Ca^{2+} → tissue

The discovery of PIEZO1 and PIEZO2, recognized by the 2021 Nobel Prize, has radically expanded our understanding of the biology of mechanosensitivity. PIEZO1 has been documented in the vascular endothelium, the smooth-muscle arterial wall, CNS vessels, and the human myometrium, where it functions as a mechanosensitive cation channel linked to the Ca^{2+}-dependent response of the tissue (Li et al., 2014; Retailleau et al., 2015; Harraz et al., 2022; Barnett et al., 2023). A key property of PIEZO1 in the vascular context is that it reads not the static state of the tissue, but rather dynamic changes of the mechanical environment, including shear stress and the associated Ca^{2+} flows (Li et al., 2014; Rode et al., 2017). This aligns well with the use of ddF as a variable that picks out short-duration episodes of maximal change of the tidal potential. Activation of PIEZO1 triggers Ca^{2+} entry into the cell and subsequently engages endothelial signaling pathways, including eNOS/NO and flow-dependent vascular regulation.

All observed wave-driven somatic signals fit this scheme:

- **CVA (stroke)** on **WAVE_event** — acute vascular decompensation at the moment of rapid hemodynamic change.
- **HYPTN (hypertensive crisis)** on **quadrature** — elevated vascular wall tone against the background of impulsive loading.
- **OBMIS (threatened miscarriage)** ↓ on **WAVE_before_quadrature** — stably across both epochs. Possible mechanism: activation of mechanosensitive circuits in the myometrium and cervix transfers part of the unstable states from a threat regime into a resolution regime.
- **OBLAB (labor)** ↑ on **FirstQ** — labor activity as a continuation of the same mechanism in the next phase.
- **INBLED (internal GI bleeding)** ↑ on **WAVE_after_quadrature** — post-impulse vascular decompensation in the wall of the gastrointestinal tract.

These tissues are united by high mechanosensitivity and the involvement of PIEZO-dependent circuits, and so their phase-coherent reaction to **ddF** appears not random but systemic.

5.4. The neuroendocrine axis and the circadian pacemaker

The suprachiasmatic nucleus (SCN) of the hypothalamus — the principal circadian pacemaker — operates as a Ca^{2+} oscillator: intracellular cAMP/Ca^{2+} signaling is not merely regulated by molecular clocks but is a necessary component of them (O'Neill et al., 2012). This makes the SCN a convergence point for signals that alter the Ca^{2+} tone of neurons and, through it, for the modulation of melatonin, cortisol, and hypothalamic–pituitary rhythms.

In the context of the present data:

- the decrease in acute cardiac arrests (**CARD/ARREST** ↓) on the full moon is consistent with lunar-dependent shifts of the neuroendocrine and behavioral background, including changes in melatonin, cortisol, and sleep parameters (Dergaa et al., 2021; Casiraghi et al., 2021).
- The increase in aggression (**AGGRESSION_TOTAL** ↑) on the new moon, reproduced in both epochs, points to a different neuroendocrine phase of the cycle, not reducible to the direct action of lunar light;
- the rise in respiratory call-outs (**RESPIR**) and syncopes (**UNC**) on the new moon looks like the autonomic–vascular profile of the same phase. Thus, the observed effects are best consistent with a model in which the lunar cycle acts through several converging channels, with the SCN serving as a node of their integration.

5.5. The behavioral layer: consciousness as the intermediate link

A principal question that cannot be sidestepped in a work devoted to the influence of the lunar cycle on criminal and medical statistics: **what exactly happens between the astronomical trigger and the registered behavioral outcome?**

The very framing of the question — about the connection between a gravitational background and crime or ambulance calls — risks sounding reductionist if the causal chain is cut at the level of "PIEZO activated → property crime decreased." This is a methodologically incorrect reduction. The actual chain contains an obligatory **intermediate link — the state of consciousness**, the emotional regulation, and the neuropsychological background of the subject.

The chain assumed in the present work:

1. **Astronomical trigger** — a ddF impulse (or a static phase) — acts through the channels described in §5.2 and §5.3.
2. **Biochemical shift** — a change in Ca^{2+} tone, in the melatonin profile, and in the activity of the HPA/HPG axes.
3. **Shift in the state of consciousness and emotional regulation** — a change in the threshold of inhibition, level of anxiety, impulsivity, libido, and risk perception. This is the **central intermediate link**, without which no connection between gravitation and behavior would exist.
4. **Behavioral output** — realized only in the presence of predisposing factors (motivation to commit a crime, a conflict, a physiological predisposition to decompensation). The lunar signal **does not create** the motive; it shifts the threshold.

This model is a classical *trigger* × *predisposition* model in population behavioral epidemiology. The small magnitude of the observed effect ($g \approx 0.03$–0.4) is fully consistent with such a model: the trigger does not act uniformly on the entire population, but rather shifts the risk for those who are already in the "gray zone." In a metropolis, where thousands

of property crimes are recorded daily, even a small shift in the threshold of self-control among the population in the "gray zone" translates into a statistically robust decrease or increase in the counts.

The state of consciousness in this model is treated **not as an abstract but as a neuropsychological** category. It is described through objectively measurable parameters: sleep quality (the amount of slow-wave phase), the level of anxiety on standard scales, the reactivity of the autonomic nervous system, attentional concentration, impulse control (Stop tests, Iowa Gambling Task), libido level, aggressive readiness. Each of these parameters is measurable, and each of them can shift under the influence of well-documented neuroendocrine changes (melatonin, cortisol, testosterone).

Direct verification of this chain at the level of individual neuroimaging and psychometrics is a natural direction for further research. Indirect verification, however, is already present in the data themselves: the observed syzygy/quadrature dichotomy in psychiatric calls (EDP ↓ globally on quadrature, EDP ↑ in public transit) shows that **the scene in which the behavioral output is realized depends on the state of consciousness**, rather than only on external factors. Before quadrature, patients with a predisposition to decompensation **do not call and do not attract attention indoors** — yet decompensate where the scene itself imposes additional loading (rhythmic vibration of transit, enclosure, sensory overload).

This same point also explains the key structural result — the decrease in property crime before quadrature (TIER1). Planned property crimes require preserved impulse control, concentration, and risk assessment. If, before quadrature, the population background shifts toward "flattening", apathy, and reduced motivation, then rational criminal activity is postponed. The signals support this interpretation: simultaneously with PROPERTY_CRIME ↓ before quadrature, EDP psychiatric calls indoors, knife conflicts, and trespass incidents all decline. The point is not that "gravity stopped the thief" — this is **a synchronous reduction of behavioral readiness** across all forms of active public activity.

5.6. The syzygy / quadrature dichotomy as a structural fingerprint

The strongest structural result of the present work is the robust **opposition of signs** between syzygy and quadrature windows across several independent physiological axes:

Axis	Syzygy (New / Full)	Quadrature (First Q / Last Q)
Property crime	↑	↓
Acute cardiac arrests (ARREST/CARD)	↓	—
Acute vascular events (HYPTN, CVA, INBLED)	—	↑
Indoor activity (burglaries, premises entries, labor)	—	↑

Axis	Syzygy (New / Full)	Quadrature (First Q / Last Q)
Street aggression	↑	—
Psychiatric calls (EDP, indoors)	—	↓

This structural bipolarity cannot be attributed to random fluctuation: it reproduces across several independent classes of outcomes and across two temporal epochs. The proposed explanation is **the different configuration of the total tidal potential**: in syzygy, the lunar and solar components are summed (maximum of the modulus); in quadrature, they partially cancel (minimum of the modulus). Biological systems respond not to "force as such" but to the sign and dynamics of the change, and **passage through the minimum of the potential in the quadrature window** creates a qualitatively different regime than passage through the maximum in syzygy.

Within the proposed model:

- In **syzygy** windows, the synchronized amplification of the phase channels dominates; the neuroendocrine system finds itself in a "pumping" regime followed by discharge — hence the behavioral activation (aggression, property and street conflicts), while the full moon can be accompanied by a **paradoxical** decrease of acute cardiac events as part of an overall circadian–autonomic shift.
- In **quadrature** windows, the biomechanical ddF impulse dominates at a reduced static potential; this is the scene in which **PIEZO-dependent tissue signals** are loudest, while behavioral activity is suppressed.

This dichotomy is the most striking structural confirmation of the **independence** of the two signal layers (phase and wave) and at the same time of their **connectedness** through the common astronomical root.

5.7. Nodal and apsidal modulators

In addition to the synodic (29.5 d) and anomalistic (27.55 d) cycles, the lunar orbit contains two long-period modulations: the **nodal cycle of 18.61 years** (precession of the orbital plane, determining the maximum lunar declination from ±18.5° at the minor standstill to ±28.5° at the major standstill) and the **apsidal cycle of 8.85 years** (precession of the line of apsides).

An observation obtained after the analysis of the results of the first run, during the calibration stage of the ddF-detector development: **the accuracy of the alignment of the registered impulsive events with the calculated lunar phase depends on the nodal phase**. In years of a high nodal parameter (2001–2005, 2024–2026 — major lunar standstill periods) the signal is detected more cleanly and is centered relative to the phase; in years of a low nodal parameter (the mid-2010s, the minor standstill period) the signal shifts and less frequently falls within the expected window. This is consistent with the physics: at a larger

amplitude of lunar declination, the vertical projection of the gravitational vector at temperate latitudes varies more strongly, which amplifies the differential loading on tissues.

The apsidal cycle (8.85 years) presumably operates in antiphase with the nodal cycle as a modulator of the location of the perigee relative to the syzygies: when the perigee systematically coincides with a syzygy, maximum amplification is observed; when it moves away, partial compensation occurs. This question belongs to the area of further investigation and requires extending the temporal window of analysis to several full nodal cycles (~40 years of data). Within the present work, we confine ourselves to fixing the fact that **the presence of long-period modulators is recommended to be taken into account when replicating SGR results in different years**, especially near the minor standstill, where the expected amplitude of the observed effects will be lower.

---

6. Safeguards against artifacts and alternative explanations

For a study claiming the existence of a connection between an astronomical parameter and socio-medical statistics, a systematic analysis of alternative explanations is mandatory. Below are the main classes of possible artifacts and the counter-arguments.

(A) Moonlight as a trivial confounder.
Objection: "the observed effects may be explained by the illuminance of the lunar disk, rather than by a gravitational signal." Counter-argument: in a metropolis, moonlight is far too small relative to artificial night lighting to explain the observed effects. Moreover, **quadrature** is physically less luminous than the full moon, yet it is precisely in these windows that the majority of vascular and wave-driven signals are concentrated; consequently, the **syzygy / quadrature** dichotomy is not reducible to illuminance.

(B) A calendar-structure artifact.
Objection: "the lunar cycle may correlate with days of the week, holidays, and seasonality." Counter-argument: all these variables are included in the OLS detrending (§3.3) — a polynomial trend of degree 4, an annual Fourier basis with 3 harmonics, day-of-week dummies, and dummies for U.S. federal holidays with ±1-day lags. The signal is estimated on residuals from which the entire calendar structure has already been removed.

(C) A COVID-period artifact.
The 2020–2021 COVID period is excluded from the main analysis as a period of extreme regime change. All TIER1 and TIER2 results are obtained on the PRE and PRE2014 epochs, which do not include COVID.

(D) A multiple-comparisons artifact.
A BH-FDR correction is applied within strata (§3.8). The TIER1 result PROPERTY_CRIME × WAVE_before_quadrature × PRE2014 × WIN3 at $p = 3.93 \times 10^{-3}$, $k = 21$ across 5 cities

remains significant after any reasonable correction. An additional safeguard is the replication in the PRE epoch (TIER2) and in pure samples (PURE_ROBBERY and others).

(E) A family-aggregation artifact.
Objection: "pooling heterogeneous codes into families may artificially create correlations." Counter-argument: pure samples (single codes) confirm the sign of the aggregated signals with $I^2 = 0\%$. Family aggregation increases power, but does not create the effect.

(F) A standard-error artifact at small *k*.
Objection: "a random-effects meta-analysis with $k = 5$ may anti-conservatively underestimate the standard error." Counter-argument: the **HKSJ correction** is applied — the modern standard for small k, explicitly conservative relative to the standard DerSimonian–Laird estimator. Applying HKSJ moves many signals down one step on the TIER scale; the results that survive HKSJ are methodologically the most reliable.

(G) An artifact of joint movement of seasonal series.
Objection: "the lunar signal and a social series may coincide in seasonal phase by chance, especially given a synodic period close to a month." Counter-argument: the synodic cycle of 29.53 d is fundamentally not synchronous with the Gregorian calendar and, over a long period, covers all possible configurations of "day of week × day of month × season." Over a five-meter sample of 5 to 25 full years of observation for each city, the phase shift between the lunar and calendar cycles fully "rolls through" all possible configurations. An additional safeguard is the circular-shift permutation in the single-database pipeline (§3.7), which explicitly tests precisely this hypothesis: even a minimum shift of ≥ 7 days destroys the observed effect, which demonstrates **phase specificity**, not a seasonal confounder.

(H) An artifact of changes in registration policy.
Objection: "the decriminalization of certain offenses in the U.S. after 2014–2015 may have distorted the temporal replication." Counter-argument: it is for exactly this reason that the PRE2014 period (through 2014 inclusive) was chosen as the principal replication sub-cohort — preceding the main waves of decriminalization (California's Proposition 47, November 2014; classification reforms in several states in 2015–2017). The fact that the TIER1 signal PROPERTY_CRIME × WAVE_before_quadrature is obtained specifically in PRE2014, and the TIER2 replication — in the full PRE epoch, **rules out** the possibility of explaining the signal by post-2014 reforms.

(I) A "winner's curse" artifact in lag selection.
Objection: "estimating the effect at the best of several lags is inflated due to optimization." Counter-argument: the headline estimate uses a **fixed** three-day window WIN3, which unites the lags {−1, 0, +1} without optimization; individual lags are reported as a diagnostic context rather than as principal results. This is a deliberate methodological decision specifically aimed against winner's curse.

(J) An artifact of researcher degrees of freedom.
Objection: "the procedures of detrending, winsorization, and gatekeeping contain free

parameters that may have been tuned." Counter-argument: all thresholds (`THR_BASE = 0.035` , `THR_STEP = 0.020` , `LOOKBACK_START = 4` , `COOLDOWN = 2` , winsorization 1%/1%, gatekeeping `mean ≥ 8/d` , `n_target ≥ 15` , `n_ctrl ≥ 60` , the TIER thresholds) **are fixed in the published code prior to the moment of the final run** and were not varied with the aim of improving the results.

In sum: none of the alternative explanations considered is capable, individually or in combination, of explaining the totality of the observed results — the TIER1 signal with I^2 = 0% in 5 independent cities, its temporal replication in two epochs, the consistency of the signs of pure samples with the family aggregation, and the syzygy/quadrature dichotomy across several independent physiological axes. This is the substantive argument in favor of the reality of the effect.

---

7. Limitations

The present study, for all the rigor of its methodology, possesses a number of principal limitations that must be explicitly stated for the correct interpretation of the results and for the planning of subsequent work.

7.1. Daily temporal resolution. All databases are aggregated at the level of the calendar day. This precludes the direct evaluation of intra-day shifts and does not allow the differentiation of, for example, night-time and day-time subsets of events. Meanwhile, the biologically expected effects of the lunar signal (especially the non-photic ones — melatonin, sympathetic–parasympathetic balance) have a pronounced intra-day structure. The use of local noon as the reference point of the astronomical calculation partially compensates for, but does not eliminate, this limitation. Further studies with hourly stratification (where the source data permit) are anticipated as a natural extension of the method.

7.2. The lack of accounting for the geomagnetic background (K-index). Geomagnetic activity was not included in the model as a covariate, although it is potentially important first and foremost for the electromagnetic and circadian channels. An elevated geomagnetic background may independently modulate the melatonin and autonomic profile and thereby be confounded with the effects attributed to channels **B** and **D**, especially during periods of elevated magnetospheric disturbance. Going forward, the inclusion of the **K-index**, or the analysis of a subset of days with quiet geomagnetic background (**K ≤ 2**), may improve the purity of the estimation of lunar-specific effects.

7.3. The lack of scene localization. In most databases, the scene in which an event is realized ("indoors / on the street / in transit") is not delineated with systematic quality. Where such stratification is possible (NYC 911, NYC EMS), it has been carried out and has yielded a structural result — a scene-dependent dichotomy (EDP). On fully aggregated criminal samples, the scene is not separated.

7.4. Population, rather than individual, level of effect. The observed shifts in event counts (~3–6% above background for property crimes, ~10–20% for individual medical indicators) are **population statistics**, not individual probabilities. For a specific individual, the lunar signal is one of many factors whose amplitude is substantially smaller than the influence of nutrition, sleep, chronic stress, and social and psychological factors. The present work demonstrates a population-level registered shift and thereby provides a basis for predictive models, yet it does not in itself validate individual prognosis without the addition of personal covariates.

7.5. Limitations of geographical extrapolation. Of the five cities analyzed, four are located directly on the ocean coast (NYC, LA, SF, Philadelphia via the Delaware Bay), and one is located on a large lake basin (Chicago, Lake Michigan). The regional magnetohydrodynamic channel (see §5.2, Channel E) is present in all five, although with different amplitudes. The extrapolation of the results to deeply continental regions, separated by 500+ km from large bodies of water, requires separate verification; the expected amplitude of the wave signal there may be lower due to the attenuation of Channel E.

7.6. Temporal localization of the epochs. The PRE2014 period includes phases of both high and low nodal parameter. However, the bulk of PRE (the extension to 2020) partially falls on the minimum of the nodal cycle (2015). This produces a heterogeneity in the expected amplitude of the effect across epochs and may partially explain why PRE2014 yields a cleaner TIER1 result than the extended PRE. A targeted analysis of the effect of the nodal phase on the reproducibility of the signal is a direction for further investigation.

7.7. Limitations of causal interpretation. The present work establishes robust **associative** links between astronomical triggers and registered events. A correct causal-inference interpretation presupposes additional verifications at the molecular level (neuroendocrine markers, laboratory models of mechanoreception in the ddF windows) and within the framework of prospective protocols with individual monitoring.

---

8. Significance and directions of practical application

If the effects described are confirmed in further independent replications — and the methodology, data, and code of the present work have deliberately been made fully open in order to facilitate such replication — a number of practically meaningful directions open up.

8.1. Forecasting the load on emergency services. The TIER1 signal of OTHER EMS (the non-specific load on the ambulance service) in the WAVE_before_quadrature window, reproduced across two epochs, indicates the potential for a deterministic component in the call-out schedule. Even a modest (5–10%) advance prediction of the daily load 1–3 days ahead can optimize the on-duty schedule, the placement of crews, and the allocation of intake beds. A similar potential exists for the 911 service in the syzygy and after_syzygy windows, in which street activity increases.

8.2. Prevention of acute vascular events. The consistent increase of HYPTN (hypertensive crisis) on quadrature and INBLED (internal GI bleeding) in the after_quadrature window indicates time windows of elevated risk for patients with the corresponding chronic conditions. In applied terms, this provides a basis for **chronopharmacological protocols** — an individualized regimen of antihypertensive therapy, anticoagulant loading, and monitoring of high-risk patients precisely in the ddF windows. Likewise — for the risk of acute cerebrovascular accident (CVA), for which the signal is registered specifically on WAVE_event.

8.3. Obstetric prevention and planning. The signals OBLAB (labor) ↑ on FirstQ and OBMIS (threatened miscarriage) ↓ on WAVE_before_quadrature are robust and reproducible. This is the basis for the precise planning of obstetric load and for understanding the "quiet" windows in which unstable pregnancies more readily resolve naturally. The present work does not draw a direct clinical conclusion; the formulation of clinical recommendations requires prospective controlled studies.

8.4. Forecasting psychiatric decompensations and panic states. The scene-dependent dichotomy of EDP (a decrease indoors, an increase in public transit on WAVE_before_quadrature) provides a basis for understanding *where* and *when* psychiatric risk concentrates. For urban mental-health services, this is a potential instrument of targeted prevention — strengthening presence at transit hubs in the wave-impulse window before quadrature; for individual patients with a diagnosis of panic disorders, it is a rationale for applying anxiety scales and behavioral strategies in elevated-risk windows.

8.5. Criminological planning. The reproducible reduction of property crime on WAVE_before_quadrature and the increase of aggression on the new moon are statistical signals at the population level that can be incorporated into models of patrol-load forecasting. It is emphasized: the application of such models presupposes **resource optimization only** (the distribution of patrols, hot-window analytics), rather than individual prediction or predictive policing.

8.6. Chronobiological personal planning. At the individual level, knowing the lunar context of one's own state — particularly in the presence of heightened nervous-system sensitivity, diagnosed autonomic dysfunction, migraine-like episodes, or bipolar affective disorder — provides a tool of **informed self-attitude**: understanding that "this is a biological window, not a catastrophe" is itself of therapeutic value within the cognitive-behavioral paradigm.

8.7. Opening a new class of research tasks. The principal methodological consequence of the present work is the introduction of **the second derivative of the tidal potential** as an autonomous variable in population epidemiology and chronobiology. The same approach, applied to other classes of outcomes (cross-country birth activity, psychiatric service contacts, the incidence of atrial fibrillation, sports traumatology, road traumatology across countries with different lighting density), may reveal previously undescribed structural regularities.

9. Data, code availability and reproducibility

9.0. Ethics, funding and competing interests.
All datasets used in this study are publicly available and fully de-identified at source by the respective city open-data programs; no human subjects were contacted and no individual-level data were obtained. The author declares no competing interests. This work is self-funded independent research with no external financial support.

In line with the principles of open science, all components of the research have been made fully available:

9.1. Source data. All seven databases used are public and accessible through the official open-data portals (for the exact URLs see Table 1, §2.1, as well as `data/raw/README.md` in the repository).

9.2. Processed daily tables. All intermediate daily CSVs (after unification, aggregation, and the application of SEMANTIC_MAP) are published in the Zenodo snapshot of the release. The volume is about 8 MB, which makes the reproduction of all results, starting from the meta-analysis step, possible on a standard personal computer.

9.3. Lunar tidal-wave files. Full daily series `waves_*.csv` with the computed values of `F`, `dF/dt`, `d²F/dt²`, lunar phase, declination, perigee proximity, and all binary triggers (`wave_*`, `phase_*`) for each city over the full covered period are included in the repository.

9.4. Analytical output files. The complete tables of correlations (`outputs/correlations/*.csv`), the meta-analysis (`outputs/verdicts/sgr_meta_family.csv`, `sgr_meta_crimetype.csv`, `sgr_forest_data.csv`), and the single-database verifications (`sgr_single_verdict_signals.csv`, `sgr_*_single_results.csv`) are published in full.

9.5. Source code. The full implementation in Python 3.10.11:

- `base.py` — processing of the "raw" databases, aggregation by local calendar days, mapping of crime and medical-call codes, removal of duplicates;
- `gen.py` — the astronomical engine, computation of F, dF, d²F, the ddF-detector of impulsive waves, generation of phase and wave triggers;
- `correlator.py` — detrending (Branch A / Branch B), computation of Hedges' g, variance, the U-test, Δ%, BH-FDR;
- `verdict.py` — random-effects meta-analysis with HKSJ correction, TIER classification;
- `single_correlator.py` — the single-database pipeline with the robustness battery (block-permutation, circular-shift, bootstrap, slice stability, placebo, Quasi-Poisson GLM);
- `single_verdict.py` — TIER classification of single-database results.

All scripts contain full comments, fixed thresholds (without hidden researcher degrees of freedom), and can be executed sequentially for the full reproduction of all results of the present work.

9.6. Licensing and citation.

- Code: MIT License.
- Processed data: Creative Commons CC-BY 4.0.
- Citation: Tysiatskii, A. (2026). *Seleno-Gravitational Rhythm (SGR): Impulsive lunar tidal waves as a modulator of human consciousness, behavior, and somatic health. A cross-city meta-analysis of criminal and medical statistics in U.S. metropolitan areas (2001–2025).* Zenodo. `[DOI: 10.5281/zenodo.20518660]`

9.7. GitHub repository: https://github.com/TemaTys/seleno-gravitational-rhythm
9.8. Permanent archive: `[DOI: 10.5281/zenodo.20518660]`

Acknowledgments

The author thanks the open-data programs of the City of Chicago, City of New York, City of Los Angeles, City of Philadelphia, and City and County of San Francisco for making the underlying datasets publicly available, and the developers of the Skyfield library and the JPL Solar System Dynamics group for the DE421 ephemeris.

References

1. Cajochen C., Altanay-Ekici S., Münch M., Frey S., Knoblauch V., Wirz-Justice A. (2013). Evidence that the lunar cycle influences human sleep. *Current Biology*, 23(15), 1485–1488.
2. Wehr T. A. (2018). Bipolar mood cycles and lunar tidal cycles. *Molecular Psychiatry*, 23(4), 923-931.
3. Lewis A. H., Grandl J. (2015). Mechanical sensitivity of Piezo1 ion channels can be tuned by cellular membrane tension. *eLife*, 4, e12088.
4. Wang S., Chennupati R., Kaur H., Iring A., Wettschureck N., Offermanns S. (2016). Endothelial cation channel PIEZO1 controls blood pressure by mediating flow-induced ATP release. *The Journal of Clinical Investigation*, 126(12), 4527–4536.
5. Wang Y., Li C., Zhang Y., Wang H., Zhu X., Li S., Sun X. (2024). PIEZO1 Promotes the Migration of Endothelial Cells via Enhancing CXCR4 Expression under Simulated Microgravity. *International Journal of Molecular Sciences*, 25(13), 7254.

6. Kyba C. C. M., Mohar A., Posch T. (2017). How bright is moonlight? Astronomy & Geophysics, 58(1), 1.31–1.32.
7. The Nobel Prize. (2021). Press release: The Nobel Prize in Physiology or Medicine 2021. NobelPrize.org.
8. Coste B., Mathur J., Schmidt M., Earley T. J., Ranade S., Petrus M. J., Dubin A. E., Patapoutian A. (2010). Piezo1 and Piezo2 are essential components of distinct mechanically activated cation channels. *Science*, 330(6000), 55–60.
9. Li J., Hou B., Tumova S., Muraki K., Bruns A., Ludlow M. J. et al. (2014). Piezo1 integration of vascular architecture with physiological force. *Nature*, 515(7526), 279–282.
10. Rode B., Shi J., Endesh N., Drinkhill M. J., Webster P. J., Lotteau S. J. et al. (2017). Piezo1 channels sense whole body physical activity to reset cardiovascular homeostasis and enhance performance. *Nature Communications*, 8, 350.
11. Retailleau K., Duprat F., Arhatte M., Ranade S. S., Peyronnet R., Martins J. R. et al. (2015). Piezo1 in smooth muscle cells is involved in hypertension-dependent arterial remodeling. *Cell Reports*, 13(6), 1161–1171.
12. Harraz O. F., Klug N. R., Senatore A. J., Hill-Eubanks D. C., Nelson M. T. (2022). Piezo1 Is a Mechanosensor Channel in Central Nervous System Capillaries. Circulation Research, 130(10), 1531–1546.
13. Barnett S. D., Asif H., Buxton I. L. O. (2023). Novel identification and modulation of the mechanosensitive Piezo1 channel in human myometrium. The Journal of Physiology, 601(9), 1675–1690.
14. O'Neill J. S., Reddy A. B. (2012). The essential role of cAMP/Ca2+ signalling in mammalian circadian timekeeping. Biochemical Society Transactions, 40(1), 44–50. DOI: 10.1042/BST20110691
15. Lucas R. J., Peirson S. N., Berson D. M., Brown T. M., Cooper H. M., Czeisler C. A. et al. (2014). Measuring and using light in the melanopsin age. *Trends in Neurosciences*, 37(1), 1–9.
16. Casiraghi L., Spies I., Chelivasi C., Fernandez-Duque E., Valeggia C., Golombek D. A. (2021). Moonstruck sleep: Synchronization of human sleep with the moon cycle under field conditions. *Science Advances*, 7(5), eabe0465.
17. Hartung J., Knapp G. (2001). A refined method for the meta-analysis of controlled clinical trials with binary outcome. *Statistics in Medicine*, 20(24), 3875–3889.
18. Sidik K., Jonkman J. N. (2002). A simple confidence interval for meta-analysis. *Statistics in Medicine*, 21(21), 3153–3159.
19. IntHout J., Ioannidis J. P. A., Borm G. F. (2014). The Hartung-Knapp-Sidik-Jonkman method for random effects meta-analysis is straightforward and considerably outperforms the standard DerSimonian-Laird method. *BMC Medical Research Methodology*, 14, 25.
20. Higgins J. P. T., Thompson S. G. (2002). Quantifying heterogeneity in a meta-analysis. *Statistics in Medicine*, 21(11), 1539–1558.
21. Benjamini Y., Hochberg Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B*

(Methodological), 57(1), 289–300.

22. Hedges L. V. (1981). Distribution theory for Glass's estimator of effect size and related estimators. *Journal of Educational Statistics*, 6(2), 107–128.
23. Sabaka T. J., Tyler R. H., Olsen N. (2016). Extracting ocean-generated tidal magnetic signals from Swarm data through satellite gradiometry. *Geophysical Research Letters*, 43(7), 3237–3245.
24. Grayver A. V., Olsen N. (2019). The magnetic signatures of the M2, N2, and O1 oceanic tides observed in Swarm and CHAMP satellite magnetic data. *Geophysical Research Letters*, 46(8), 4230–4238.
25. Hocke K. (2025). Modulation of the lunar semidiurnal tide in GNSS TEC by the variable Earth-Moon distance. *Frontiers in Astronomy and Space Sciences*, 12, 1585247.
26. Weichenberger M., Bauer M., Kühn S., Koch C., Müller J., et al. (2017). Altered cortical and subcortical connectivity due to infrasound administered near the hearing threshold – Evidence from fMRI. PLoS One, 12(4), e0174420.
27. Rycroft M. J., Harrison R. G., Nicoll K. A., Mareev E. A. (2008). An Overview of Earth's Global Electric Circuit and Atmospheric Conductivity. Space Science Reviews, 137, 83–105.
28. Bevington M. (2015). Lunar biological effects and the magnetosphere. Pathophysiology, 22(4), 211–222.
29. Rhodes B. (2019). Skyfield: High precision research-grade positions for planets and Earth satellites. Astrophysics Source Code Library, ascl:1907.024.
30. Folkner W. M., Williams J. G., Boggs D. H. (2009). The Planetary and Lunar Ephemeris DE 421. IPN Progress Report, 42-178.
31. Dergaa I., Romdhani M., Fessi M. S., Ben Saad H., et al. (2021). Does lunar cycle affect biological parameters in young healthy men?. Chronobiology International, 38(6), 933–940.
32. Tessmar-Raible K., Raible F., Arboleda E. (2011). Another place, another timer: Marine species and the rhythms of life. *BioEssays*, 33(3), 165–172.
33. Wang C., White-Schwoch T., Strüfeld T., Connor C. T., Kirschvink J. L. (2019). Transduction of the geomagnetic field as evidenced from alpha-band activity in the human brain. *eNeuro*, 6(2), ENEURO.0483-18.2019.