World Tension Meter

Geopolitical Stress & Inflation Pressure Index

Methodology Whitepaper

Version 3 — Engine v3.5 · July 2026

How a fully automated instrument measures the world's tension
twice a day, from free public data, with every rule in the open.

Supersedes the v1 whitepaper (February 2026, first-generation engine).

worldtensionmeter.com · Dataset: CC BY 4.0 · contact@worldtensionmeter.com

1 · What the World Tension Meter is

The World Tension Meter (WTM) is a fully automated geopolitical risk index. It compresses armed conflict, energy stress, trade disruption, financial pressure and global media coverage into one daily 0–100 score, published on a fixed schedule with no human editing: full analysis twice a day (00:10 and 12:10 UTC), market data every ten minutes. The design goal is an **instrument, not an opinion** — every number on the site is produced by the engine described in this paper, from named public sources, under rules documented here.

2 · Design principles

Honesty over availability

If a data source fails, the engine reuses the last known good value and publicly flags it “estimated” — in the JSON (_meta.estimated_fields) and on the site. Nothing is ever invented. If data stops flowing entirely, the interface says “Delayed” rather than quietly showing stale numbers.

A permanent, append-only record

Every daily score is appended to a public ledger and never edited afterward. The record began 2026-03-23 and grows by exactly one row per day; the Backtesting page charts it raw. In time, the ledger itself becomes the strongest evidence for or against the method.

No single point of failure

The five factors come from separate data worlds, so one noisy source cannot move the composite. Every fetch has a timeout, one retry, a cached fallback and an estimated flag. All sources are free and keyless; the engine self-limits its request budget rather than trusting rate limits.

3 · The composite score

The World Tension Score is a weighted composite of five independently measured factors, each normalized to 0–100:

Factor	Weight	Primary inputs
Conflict	25%	Global conflict-coverage intensity (GDELT), event-level data where available, conflict share of world ne
Energy	20%	Brent crude, natural gas, oil volatility, gold/oil ratio
Trade	20%	Baltic Dry Index, chokepoint & tariff coverage share, food-commodity stress
Finance	20%	VIX, 10-year Treasury yield, dollar-index deviation, gold
Media	15%	GPR coverage-share signal, headline fear density, coverage tone

On top of the composite, a capped **event-shock overlay** (§5) reacts to detected crisis events, and **fear physics** (§5) governs how the final published score moves.

4 · The media signal: coverage share, not keywords

The media and conflict signals follow the approach of the academic **Geopolitical Risk (GPR) index** (Caldara & Iacoviello, Federal Reserve Board): measure what *share* of the world's news coverage is tension-related, then rank today against that same query's own history. The engine queries the GDELT Project's global monitoring volume (timeline volume as a share of all coverage) and converts today's value to a percentile of its own baseline. A reading of 85 therefore means: coverage genuinely ranks in the top 15% of the baseline window — not that a keyword list fired.

Why the baseline window is ten years

The window length was settled empirically, in public, on live data. A 12-month baseline proved war-contaminated: every month in it was wartime, so war ranked as "normal" (the observed 12-month minimum share was 8.56% — there were no calm months to anchor zero). A 5-year window (2021–2026) had the same defect in slower motion: an active major war ranked near the 56th percentile. The production window is **ten years**: the genuinely calm 2016–2019 era anchors the floor, so wartime ranks where human judgment says it should. Baselines refresh weekly, are cached against outages, and the percentile is soft-bounded (2–98) so the score can never mathematically pin to 0 or 100.

5 · Fear physics and event shocks

Real fear is asymmetric, and so is the index. When the composite rises, the published score jumps 85% of the gap immediately. When it falls, the score bleeds at most **three points per day**, anchored by a persisted daily floor. A major attack therefore spikes the meter within one cycle and holds for days — the way markets, and people, actually behave.

Independently, a shock detector scans headlines for crisis events (invasions, strikes on infrastructure, nuclear threats, market crashes). Detected events add a temporary overlay capped at **+18 points**, each decaying with its own half-life. Question headlines ("Will X happen?") are excluded. Worked example from the live record: on 2026-07-23 the engine detected "Kuwait–Iraq border crossing hit in drone attack" and applied +14 within minutes of publication — visible in the public ledger.

6 · Sub-indexes

Food Security. Wheat (45%), corn (35%) and rice (20%) futures, each ranked against multi-year price history with unit-flip protection.

Currency Stress. Dollar-index deviation from neutral (50%) plus averaged stress across EUR, JPY, CNY and GBP pairs (50%).

Shipping Stress. Baltic Dry Index level (60%) plus coverage share of chokepoint, blockade and tariff disruption (40%). When the freight feed is unavailable the last known value is reused and flagged estimated.

Sanctions Pressure. Coverage share of sanctions activity (50%), sanctioned-currency stress (30%), standing-regime baseline (20%).

Energy Security. Brent (40%), natural gas (30%), oil volatility (15%), gold/oil ratio (15%), each ranked against multi-year history.

Gold & Silver Fear. Metal levels against multi-year history (50%), VIX (30%), 10-year yield (20%).

Crypto Fear & Greed. 24h momentum (30%), 7-day momentum (25%), volatility vs 30-day norm (20%), BTC-dominance shift (25%), blended 80/20 with the independent alternative.me index as an external cross-check.

7 · Data sources and discipline

All production sources are free and keyless: the GDELT Project (global news monitoring), public market data feeds (Stooq primary, Yahoo Finance fallback), major world and business RSS wires, and alternative.me for the crypto cross-check. Event-level conflict records (ACLED) are used where available but are never load-bearing. The engine enforces its own budget: at most 8 GDELT requests per run with 15-second pacing, 48-hour value caching, and prioritized weekly baseline refreshes — designed to be a polite, sustainable consumer of public infrastructure for decades.

8 · Reading the score

0–39	**Calm**	Normal conditions. No geopolitical inflation pressure.
40–54	**Stable**	Background stress present but contained.
55–67	**Moderate**	Elevated stress across multiple signals. Monitor actively.
68–77	**High Tension**	Rising costs and supply risk. Review exposure.
78–100	**Extreme**	Multiple simultaneous crises. Expect disruption.

9 · Limitations, stated plainly

Scores earlier than the live ledger are modeled backcasts — useful for shape, not exact values. Coverage-share signals inherit the biases of global media attention. Event-level conflict granularity depends on an optional third-party source. The index is informational only and is not financial advice. Where the method changes, the change and its reasoning are published — as the baseline-window history in §4 demonstrates.
