

World Tension Meter

Methodology Whitepaper v4 — engine v4.2, September 2026

Abstract. World Tension Meter is a free, independent, open observatory that converts the world's news, public attention and open market prices into twelve daily indexes on a 0–100 scale, recomputed every ten minutes and kept forever. This paper documents engine v4: an in-house signal engine that reads a corpus of forty-plus free public news feeds, public attention data and public social signals, with no dependence on any single external provider; the nine measured indexes and three derived composites; the smoothing, shock and calibration rules; the conflict lifecycle; the monthly story layer; and the governance rules that keep the record honest. Weights and thresholds are stated exactly so that every published number can be reproduced from the method.

Published 2026-09-17 · Licence CC BY 4.0 · Cite as: World Tension Meter (2026), Methodology Whitepaper v4, worldtensionmeter.com/data/

1. Purpose and principles

Readable risk. Institutional risk products are paywalled, slow and written for specialists. The observatory publishes a number, a colour and a sentence on a page that loads in a second, and explains every month why the number was what it was.

Five principles govern every design decision. (1) *Independence of engine:* scores come from our own engine reading free public sources; when an enhancement source is unavailable the built-in method continues and the number still publishes. (2) *Free forever:* no paywall, advertising, premium tier or data sales; data and API are CC BY 4.0. (3) *Open method:* weights, zones and rules are public and change only with a versioned release recorded in the changelog. (4) *Honesty about limits:* estimated fields are flagged, earlier calibrations are noted, a value we cannot compute is shown as missing. (5) *Numbers are never written by AI:* language models classify headlines and write plain-English summaries from facts already computed; every score is arithmetic on measured signals.

2. Architecture and cadence

The engine runs every ten minutes (a "fast" run: market data, smoothing, publication) and twice a day at 00:10 and 12:10 UTC (a "full" run: corpus fetch, classification, attention data, conflict lifecycle, AI Risk, daily brief). Fulls are self-scheduled by the engine, so the cadence does not depend on the host's clock configuration. A publisher rewrites live values, meanings and generated pages into the site after every run; a nightly job writes monthly stories and timelines. Every run is dual-written to JSON (served to the site and the API) and to a relational ledger (runs, scores, signals, articles, events, conflicts, market data, monthly stories).

Source class	What it feeds	Failure behaviour
Own news corpus: 40+ free public feeds across world, business, energy, defence and AI	tension, trade and sanctions coverage shares; conflict activity; AI Risk; monthly headlines	feeds are plural; a dead feed lowers volume only
Public attention data (encyclopedic page-view counts on tension, trade, sanctions and AI-safety topics, three-year history)	attention percentiles blended into the coverage signals	optional; blend re-weights without it
Public social signals (open community news feeds)	small social component of the coverage blend	optional
Open market prices (oil, gas, metals, grains, equity volatility, yields, dollar index, crypto, freight)	the seven market-based indexes and the World Tension factors	last known value carried and flagged as estimated

Source class	What it feeds	Failure behaviour
Language-model classification (enhancement only)	headline tension / AI-risk refinement; plain-English brief and monthly narratives	keyword scoring and template text take over automatically

3. The twelve indexes

All indexes share one scale, 0–100, and one zone ladder: 0–39 Calm, 40–54 Stable, 55–67 Moderate, 68–77 High Tension, 78–100 Extreme. The Crypto index keeps its conventional Fear/Greed labels. Each index carries a plain-English "what this means" line chosen by zone.

Index	Measures	Main inputs
World Tension	the headline score: composite of five factors plus shock	conflict, energy, trade, finance, media factors
Food Security	grain prices against their own multi-year history	wheat, corn, rice
Currency Stress	dollar strength and stress across major pairs	dollar index and major pairs
Shipping Stress	freight cost plus coverage of chokepoints, blockades and tariffs	freight index, corpus coverage
Sanctions Pressure	how hard the sanctions machine is running	sanctions coverage share, sanctioned-currency stress, regimes in force
Energy Security	oil, gas, oil volatility and the gold-oil ratio	oil, gas, gold
Gold & Silver Fear	safe-haven demand	metals vs history, equity volatility, 10-year yield
Crypto Fear & Greed	crypto sentiment	momentum, volatility, dominance, independent cross-check
AI Risk	risk that AI causes serious harm or acts beyond human control	AI-risk corpus share, top-headline severity, attention
Inflation Pressure (derived)	how hard world markets push the cost of living	energy 35%, food 25%, shipping 20%, currency 20%
Supply Chain Risk (derived)	risk that goods stop moving or get expensive to move	shipping 40%, energy 25%, sanctions 20%, conflict factor 15%
Market Fear (derived)	what investors are doing with their money	Gold & Silver Fear 50%, (100 – Crypto F&G;) 50%

4. The World Tension score

4.1 From headlines to shares. Every headline in the corpus is scored for tension by keyword tiers (extreme +12, high +6, moderate +3 on a base of 20; named-actor +8; nuclear +15; oil-spike +10; clamped 5–99). A language model refines the classification in batches of 110 and its result is blended 60/40 with the keyword score; if the model is unavailable the keyword score stands alone. Three coverage shares are then computed: the share of world/defence headlines reading as tension, the share of business/world/energy headlines about trade disruption, and the share about sanctions.

4.2 Percentiles against our own history. Each share is placed as a percentile against the engine's own daily record of that share. A calibration bridge anchored the first percentiles to the previous engine's levels and fades out linearly over the first 120 days of record, so the public score showed no artificial jump at the cutover. The three coverage percentiles are blended with public-attention percentiles (own corpus 0.45, attention 0.40, social 0.15; the weights renormalise when a component is missing).

4.3 Factors and composite. Five factors are formed from those signals and market inputs and combined as a weighted mean: conflict 25%, energy 20%, trade 20%, finance 20%, media 15%. The result is clamped to 0–100.

4.4 Shock overlay. Event-level headlines (attacks, strikes, blockades, nuclear developments) add a shock of up to +18 points, decaying with a half-life so that a single loud day does not become a plateau; a guard discounts question-form headlines.

4.5 Fear physics. The published score moves up quickly and down slowly: 85% of any upward gap is applied immediately, while the score may fall at most 3 points per day from a day-anchored floor. This mirrors how risk is experienced: alarm is instant, reassurance is gradual.

5. AI Risk Index

Headlines from dedicated AI feeds and AI-related items elsewhere are scored on a risk ladder: extreme (agent escape, self-replication, autonomous weapons, AI-driven attacks), high (frontier capability jumps, major misuse, lab safety failures), moderate (regulation, policy, backlash). The index combines the share of AI headlines reading as risk events (40%), the mean severity of the eight strongest headlines (40%) and public attention on AI-safety topics as a percentile of its three-year history (20%), smoothed with an exponential moving average ($\alpha = 0.5$). The top headlines are published with the score so the number is always traceable to its evidence.

6. Derived indexes and divergence alerts

The three derived indexes are fixed weighted averages of indexes already computed (weights in the table above); they add no new source and can fail only if their components fail, in which case the remaining components renormalise when at least 60% of the weight is present. Two families of divergence alert are published: headlines versus markets, when World Tension and Market Fear differ by 20 points or more; and self-divergence, when energy, food, shipping, currency or safe-haven demand sits 15 points or more from its own 30-day average.

7. Conflict roster and lifecycle

Twenty-plus armed conflicts are tracked. Each is recognised in the corpus by its own set of terms (places, actors, forces). Per run the engine counts mentions in tension-scored headlines and, within them, violent mentions. Trend compares the last 30 days of mentions with the 30 before: escalating at 1.25×, cooling at 0.75×, otherwise active. No real activity (violent mentions ≥ 2 or mentions ≥ 5 in a run) for 30 days marks a conflict dormant, still listed but excluded from regional scoring; 120 days retires it to the archive, from which it revives automatically if activity returns. Countries outside the roster that draw four or more high-tension conflict headlines in a run on at least three of the last fourteen days are surfaced as candidates for human review; nothing is added to the roster automatically. The lifecycle clock started on 2026-09-17; no conflict is retired on a guess.

8. Monthly stories and the archive

The daily record is append-only from 2026-03-23. Once a month closes, every index receives a locked write-up: the biggest daily moves with the headlines of those days, the top tension and calming headlines, a per-index table and a plain-English narrative generated from those facts only. The current month refreshes nightly; closed months never change. Yearly archive pages and per-year CSV files are generated automatically, and a new archive year is created on 1 January.

Calibration note. Scores on or before 2026-07-23 were computed by the first-generation engine, which ran on a higher calibration. Pages and API responses covering that period carry this note; comparisons should be made within a period or through the monthly stories.

9. Governance and reproducibility

The engine is versioned; the live version is printed on every page and in every API response. Method changes ship only as a new version with a dated changelog entry. Estimated inputs are flagged per field. All data is available through the open API (v1: today, index, history, day, monthly, conflicts, status) and the data portal under CC BY 4.0. The observatory is self-funded and accepts no advertising, sponsorship, investment or payment; its independence statement is published at worldtensionmeter.com/independence/.

Limitations. Coverage-based signals measure what newsrooms report, not ground truth; they can lag in under-reported regions and over-react to media cycles, which the percentile framing and fear physics partly correct. Market inputs depend on public feeds that occasionally stall (flagged). The indexes are informational and are not investment, legal or security advice.

World Tension Meter — worldtensionmeter.com — an open observatory. Method: [/methodology/](#) · API: [/api/](#) · Data: [/data/](#) · Changelog: [/changelog/](#) · Status: [/status/](#)