METHODOLOGY NOTE

Israel–Türkiye Conflict Probability Simulator

Model architecture, pathway justification, and calibration rationale

Author: Dr. Mosab Hawarey
URL: https://mosab.hawarey.org
Social: @DrHawarey (Twitter / LinkedIn)
Version: 2.0 · companion to the interactive simulator (single-file HTML)
Online: https://mosab.hawarey.org/israel-turkiye-war-simulator.html
Date: 24 June 2026 (*updated*)
Location: Amman, Jordan
Horizon modeled: mid-2026 → 2032 (annual steps)

This document explains how the simulator is built and why each number is set where it is, so that an independent analyst can check the structure and priors against the open-source record. It is an analytical instrument, not a forecast, and makes no claim about the intentions of any government.

Contents

1. Purpose and Scope

The simulator estimates how the Israel–Türkiye relationship may evolve between mid-2026 and 2032, and in particular the probability that the relationship slides into direct interstate war within that window. It is the second in a family of geostrategic Monte Carlo tools; the first modeled the viability of the Islamabad memorandum between Iran and the United States. The two tools share a visual grammar and a Monte-Carlo-plus-decision-trellis design, but the engine here is rebuilt for a forward prediction rather than a retrospective viability assessment.

Three properties were treated as non-negotiable design requirements, given the sensitivity of the subject: **transparency** (every probability is a traceable function of named inputs), **symmetry** (neither state is privileged as aggressor or victim; either can initiate, and each carries its own escalation drivers and restraints), and **falsifiability** (the priors are stated explicitly so they can be disputed and re-tuned).

What the tool is not: it is not a forecast, not an intelligence product, and not a statement about what either government intends to do. Its outputs are conditional on the analyst-set dials and on a set of documented structural assumptions. They should be read directionally — as a disciplined way to reason about which factors move the risk and by how much — not as point predictions.

2. Model Architecture

Each simulated future is a path through six annual transitions (2027, 2028, 2029, 2030, 2031, 2032), beginning from the relationship's documented mid-2026 condition. The system occupies one of three active **momentum states** in any given year, with war as an absorbing fourth state:

- **Détente** — active deconfliction, diplomatic thaw, cooperative trajectory.
- **Managed Rivalry** — cold standoff with a functioning deconfliction channel. This is the starting state, reflecting the mid-2026 baseline.
- **Confrontation** — broken deconfliction, proxy clashes, or direct incidents short of sustained war.
- **War** — sustained direct armed conflict; an absorbing terminal state.

In each year the model first tests whether war ignites (Section 3). If it does not, the state moves up or down according to a net momentum score (Section 5). At the 2032 horizon the surviving non-war paths resolve into one of three end-states — Normalization, Cold Peace, or Frozen Confrontation (Section 6). Approximately 10,000 independent paths are simulated; the share ending in each terminal is the reported probability. The simulation uses a seeded pseudo-random generator so that a given dial configuration is reproducible, while the headline *Run* reseeds to expose Monte Carlo variation.

3. The War-Onset Mechanism: Noisy-OR Over Seven Pathways

War does not arise from a single switch. In each year it can ignite through any of seven documented pathways, each with its own small annual probability driven by the relevant dials and by the current state. The pathways are combined using a **noisy-OR** rule — the probability that at least one fires:

$$P(\text{war this year}) = 1 - \prod (1 - p_i)$$

This is preferable to summing the pathway probabilities (which can exceed one) or to a single lumped constant (which hides the mechanism). Each pathway probability is the product of a base

rate, the dials that drive it, and a state multiplier. The base rate is an analyst-assigned prior, not an empirically estimated frequency; the justification for its relative size is given pathway-by-pathway below, and the aggregate is anchored to open-source estimates in Section 8.

State multiplier (b_m). Every pathway is scaled by the current state: ×0.5 in Détente, ×1.0 in Managed Rivalry, ×2.0 in Confrontation. The premise — that a relationship already in open confrontation is materially more combustible than one in détente — is uncontroversial and is reflected throughout the escalation literature.

Capability creep. Türkiye's effective forward-deployment term rises modestly from 2028 (entry of the indigenous KAAN fighter into service) and again from 2030 (Eurofighter deliveries), reflecting a narrowing of the near-term qualitative gap over the horizon.[1]

Election-year amplifier. In the 2028 transition, the two Türkiye-driven pathways (accidental clash and in-theatre pre-emption) are amplified by Türkiye's election cycle, reflecting the documented domestic incentive for a harder posture in an electoral year.

3.1 Pathway summary

Pathway	Base rate	Primary drivers (dials)
1. Accidental clash / deconfliction failure	0.020	(1 − deconfliction) × Israeli strike tempo × Turkish deployment
2. Turkish in-theatre pre-emption	0.015	pre-emption propensity × Turkish deployment × Erdoğan posture
3. Israeli escalation	0.012	Israeli strike tempo × Israeli threat-framing
4. Eastern-Mediterranean naval incident	0.010	naval friction (Greece–Cyprus bloc, gas, flotillas)
5. Proxy war crosses threshold	0.014	Syrian proxy-war intensity
6. Third-party spoiler	0.010	spoiler activity (remnant Iran, ISIS, provocations)
7. Leadership-strike / homeland tail	0.006	decapitation-doctrine propensity (suppressed outside Confrontation)

All seven are then multiplied by the structural-restraint factor of Section 4 before being combined. Base rates are annual and pre-restraint.

3.2 Why each pathway is ranked as it is

Pathway 1 — Accidental clash / deconfliction failure (highest base rate, 0.020). The single most likely route to a direct exchange is an unintended collision in contested Syrian airspace if the Baku/US-brokered deconfliction channel fails. Israeli air operations over Syria rose sharply after the fall of the Assad government, and the March–April 2025 strikes on the T4/Tiyas and Palmyra airbases — as Turkish teams assessed them for deployment — brought the two militaries closer to contact than at any prior point. It is given the highest base rate and is driven by the interaction of deconfliction breakdown, Israeli strike tempo, and Turkish forward presence.[2]

[1] On KAAN service entry (Block 10, targeted 2028) and the 2025 Eurofighter Typhoon agreement (deliveries ~2030), see the simulator's source evidence base and contemporaneous defense reporting; Türkiye was removed from the F-35 program in 2019 over its S-400 purchase.

[2] Israeli strike tempo in Syria after December 2024, the T4/Palmyra episode of March–April 2025, and the subsequent Azerbaijan-hosted deconfliction talks are documented in the simulator's evidence base, drawing on Stimson Center analysis and contemporaneous reporting.

Pathway 2 — Turkish in-theatre pre-emption (0.015). Ankara is far likelier to strike Israeli assets inside Syria — framed as defensive force protection — than to strike Israel proper, which would forfeit NATO sympathy and US backing. The model therefore defaults the Turkish first move to the Syrian theater. At 0.015 it carries the highest base rate of the Türkiye-initiated routes, because where it occurs it is by definition a direct military engagement.[3]

Pathway 5 — Proxy war crosses threshold (0.014). A sustained proxy contest — Turkish-backed Syrian forces against Israeli-aligned Druze or Kurdish elements — is the most probable form any broader "conflict" takes, and the likeliest bridge from rivalry to direct clashes. It ranks just behind in-theatre pre-emption because it can run hot without either capital choosing war, and it converts into interstate war less directly than an air-to-air collision.[4]

Pathway 3 — Israeli escalation (0.012). A deliberate Israeli decision to strike Turkish forces, driven by strike tempo and by the "new Iran" threat-framing voiced by some Israeli figures after the 2026 Iran campaign. It is ranked below the accidental and proxy routes because a deliberate, premeditated Israeli strike on a NATO member is a higher political threshold than an unintended clash.[5]

Pathways 4 and 6 — Naval incident and third-party spoiler (0.010 each). An Eastern-Mediterranean naval incident (around the Greece–Cyprus alignment, contested gas fields, or flotillas) and a third-party spoiler (remnant Iranian networks, ISIS, or a provocation) are real but lower-probability ignition routes than the Syrian land/air theater. They are set equal and modest, and the naval pathway is dampened when the relationship is in Détente.[6]

Pathway 7 — Leadership-strike / homeland tail (lowest base rate, 0.006). A decapitation strike on Turkish leadership — the playbook some observers drew from the 2026 Iran campaign — is a tail-of-the-tail event. Türkiye is a NATO member with no nuclear program serving as a casus belli and sits behind layered defenses deep in Anatolia. It is given the lowest base rate, is suppressed by a factor of four outside the Confrontation state, and is switched off entirely when the analyst disables bolt-from-the-blue war. It is retained only because the "new Iran" rhetoric makes it a non-zero branch, not because it is plausible.[7]

4. Structural Restraint Multiplier

Every pathway probability is multiplied by a single structural-restraint factor that captures the forces suppressing escalation:

$$R = (1 - 0.35 \cdot US) \cdot (1 - 0.30 \cdot NATO) \cdot (1 - 0.15 \cdot Baku) \cdot (1 - 0.20 \cdot Iran\text{-}reconstitution)$$

- **US engagement (weight 0.35)** is the largest single brake, reflecting the consensus that American leverage over both partners is the primary restraint on the relationship.
- **NATO restraint (0.30)** captures the elevated cost of striking, or being seen to provoke, a member of the alliance.

[3] The evidence base distinguishes Turkish pre-emption against Israeli assets in Syria (more plausible) from pre-emption against Israel proper (low); only the in-theatre variant is treated as a near-term route to direct conflict.

[4] On competing Syrian clients and the summer-2025 Suwayda/Druze episode, see the evidence base and Crisis Group / Washington Institute commentary on the post-Assad order.

[5] On the "Türkiye is the new Iran" framing in Israeli public debate, see the evidence base; the model treats it as an input (threat-framing dial), not as an established policy.

[6] On the Eastern-Mediterranean dimension and the Israel–Greece–Cyprus trilateral, and on the October 2025 flotilla episode, see the evidence base.

[7] The evidence base assigns the decapitation branch a sub-1% multi-year prior and treats it as reachable essentially only from an already-active confrontation; the model mirrors that by gating it on state and on the bolt-from-the-blue toggle.

- **Azerbaijani mediation (0.15)** reflects Baku's good offices and the shared interests that channel disputes away from force.
- **Iran reconstitution (0.20)** is the de-escalatory case of the single biggest long-run swing: a reviving Iran restores the shared threat that historically aligned Ankara and Jerusalem. Its escalatory mirror — a collapsed Iran removing that common enemy — is the low default value of the dial.

The weights are ordinal judgments — US first, NATO second, then Iran trajectory, then Baku — consistent with how the cited analyses rank these brakes. They are deliberately multiplicative so that no single brake can drive risk to zero, and so that brakes compound when several are strong at once.

5. Band Dynamics and Momentum

When war does not ignite in a given year, the state moves according to a net momentum score: a weighted sum of the stabilizing dials (deconfliction, US engagement, Azerbaijani mediation, Iran reconstitution, post-2026-vote de-escalation, NATO restraint) minus the escalatory dials (Israeli strike tempo, Turkish deployment, Erdoğan posture, Israeli threat-framing, proxy intensity, naval friction, election pressure). A positive score makes a move toward Détente more likely; a negative score makes a move toward Confrontation more likely; the larger the magnitude, the stronger the pull. The 2028 transition carries an additional downward nudge from the Türkiye election cycle.

The functional form (a base 0.34 probability of improving or worsening, shifted by 0.42 × momentum) is inherited from the validated first-generation model. It is intentionally simple and saturating: extreme dial settings push the move probabilities toward certainty without ever exceeding it.

6. Terminal Resolution at the Horizon

At 2032, each surviving path resolves into Normalization, Cold Peace, or Frozen Confrontation. The split depends on the state the path occupies at the horizon and on a quality score drawn from the stabilizing dials. A path in Détente resolves overwhelmingly to Normalization; a path in Managed Rivalry resolves mostly to Cold Peace, with tails into Normalization or Frozen Confrontation; a path in Confrontation resolves mostly to Frozen Confrontation. The three probabilities are constructed to be non-negative and to sum to one for every state and every dial setting — a property verified directly (Section 9). The same resolver governs both the Monte Carlo engine and the interactive “walk the tree” mode: when an analyst walks a single path to 2032, its terminal is drawn from this identical distribution rather than assigned deterministically, so the two views of the model can never disagree on how a horizon state resolves.

7. Input Variables and Default Settings

The sixteen dials and their default values are listed below. Defaults encode the open-source mid-2026 picture; they are starting points, and the entire purpose of the tool is to let an analyst move them. “Direction” indicates whether raising the dial raises (▲) or lowers (▼) war risk.

Variable	Default	Dir.	Basis
Syria deconfliction integrity	0.62	▼	Holding but fragile; the Baku/US channel that has prevented miscalculation.
Israeli strike tempo in Syria	0.62	▲	Elevated; strike volume rose sharply after the fall of Assad.

Variable	Default	Dir.	Basis
Turkish forward deployment	0.45	▲	Partial; base-deployment plans were checked by the deconfliction channel.
Syrian proxy-war intensity	0.50	▲	Live (Druze, Kurdish/SDF frictions).
US engagement / mediation	0.60	▼	Active; mediation offered. The primary external brake.
NATO restraint weight	0.70	▼	Strong structural deterrent against striking a member.
Azerbaijani mediation	0.58	▼	Active Baku channel.
Erdoğan posture	0.68	▲	Highly adversarial rhetoric.
Israeli threat-framing of Türkiye	0.55	▲	"New Iran" framing rising but not settled policy.
Türkiye 2028 election pressure	0.40	▲	Distant in 2026; rises into 2028.
Post-2026-vote de-escalation	0.45	▼	Uncertain; a post-Netanyahu government would lower framing.
Iran reconstitution	0.30	▼	Low near-term after the 2026 campaign; the escalatory case.
Eastern-Med naval friction	0.45	▲	Moderate; Greece–Cyprus alignment deepening.
Third-party spoiler	0.40	▲	Background risk (remnant Iran, ISIS).
Leadership-strike doctrine	0.20	▲	Very low; tail branch.
Pre-emptive-strike propensity	0.30	▲	In-theatre pre-emption plausible; homeland pre-emption low.

Source for the default settings: the simulator's evidence base, which compiles the 2024–2026 record on the relationship, the Syrian theater, the military balance, alliance constraints, and domestic drivers.[8]

8. Calibration and Anchoring

Because the pathway base rates are priors rather than measured frequencies, the aggregate output is anchored to the public record in two ways.

Market anchor. A prediction market priced roughly 14% for any direct Israel–Türkiye military clash before 2027 — and an equivalent 2025 market resolved at zero. Crucially, that figure covers any use of force, a far lower bar than war. The model is tuned so that its implied probability of a limited direct clash is broadly consistent with that order of magnitude, while the probability of sustained war sits well below it.[9]

Consensus anchor. The qualitative consensus across the cited institutions is that direct war is unlikely but the underlying rivalry is hardening and the danger of miscalculation in Syria is real. The default output is calibrated to express exactly that: a dominant mass on Cold Peace and the sub-war outcomes, with war as a low-double-digit cumulative tail over the six-year horizon.

Resulting default output. At the default settings the model returns approximately: Normalization 23%, Cold Peace 37%, Frozen Confrontation 28%, and War 12% cumulative by 2032 (roughly 2% per year, rising as capability and election effects accrue). Pushing all dials to their most

[8] Underlying open-source inputs are drawn from think-tank and press reporting (Brookings, Crisis Group, Stimson Center, the Washington Institute), the SIPRI Yearbook 2026 for the nuclear estimate, and force-comparison compilations whose figures vary by source and are therefore used as ranges, not point values.

[9] Prediction-market figures are moderate-volume and measure any clash, not war; their narrative summaries are partly machine-generated. They are used as a directional anchor only. No think tank or commercial risk firm publishes a competing numeric Israel–Türkiye war probability.

dangerous settings drives cumulative war into the high-sixties percent; pushing them all to their safest drives it to a fraction of a percent. These bounds are reported so the reader can see the full range the structure permits.

9. Validation

The engine was tested directly (not merely inspected). The following checks all pass:

- The per-year war probability is a valid probability in [0,1] across 300,000 randomized extreme dial combinations — no out-of-range values, no NaN.
- Only valid end-states are ever produced; terminal counts sum exactly to the number of simulated paths (no lost or double-counted futures).
- War transitions reconstructed from the path edges sum exactly to the War total — internal accounting is consistent.
- War risk is monotonic in the master variable: it never rises as the deconfliction dial is increased.
- War hazard increases strictly with the state (Détente < Managed Rivalry < Confrontation), as designed.
- All sixteen dials move war risk in their intended direction — verified exhaustively, confirming there is no asymmetric or inverted wiring.
- The terminal resolver's three outcome probabilities are non-negative and sum to one for every state and dial setting.
- All four outcomes — including War — are reachable in the interactive “walk the tree” mode as well as in the Monte Carlo run, and walk mode resolves the horizon through the same probabilistic resolver as the Monte Carlo engine (Section 6), so a single walked path is a faithful draw from the same terminal distribution rather than a deterministic shortcut.

The sensitivity panel in the tool reports, for each variable, the change in cumulative war probability from a one-notch move toward danger, holding all else fixed. Under the defaults the largest movers are deconfliction integrity, Israeli strike tempo, and US engagement — the variables the evidence base identifies as decisive.

10. Limitations, Uncertainty, and Neutrality

Priors, not measurements. The seven base rates and the restraint weights are reasoned judgments calibrated to open sources. They are the most contestable part of the model and are stated explicitly so they can be challenged. Reasonable analysts may set them differently; the tool is designed for exactly that.

Compounding uncertainty over time. The further past 2026 the model reaches, the wider the true error bars — particularly around Iran's trajectory and the posture of external powers. Cumulative figures should be read with that widening in mind.

Contested underlying data. Force balances, strike counts, and casualty figures vary by source and are sometimes partisan. Where they enter the model they are treated as ranges and as relative magnitudes, not precise values.

Structure, not intelligence. The model encodes documented structure — sequencing, theaters, brakes, and ignition routes. It has no access to classified information and infers nothing about either side's actual decision-making.

Neutrality. Both states are modeled symmetrically. Either can initiate; each has its own escalation drivers and its own restraints; the assessment panel presents an Israel lens, a Türkiye lens, and a structural lens on equal terms. No pathway, label, or default encodes a judgment that one party is the aggressor. This is a tool for thinking about risk, offered in the hope of clarifying rather than inflaming a dangerous rivalry.

Source Base

The model's inputs are drawn from a compiled open-source evidence base covering the Israel–Türkiye relationship from 2024 to 2026. Principal categories of sources:

- Think-tank and policy analysis: Brookings, the International Crisis Group, the Stimson Center, and the Washington Institute, among others.
- Quantitative references: the SIPRI Yearbook 2026 (nuclear estimate) and standard force-comparison compilations (used as ranges).
- Market signal: publicly traded prediction-market pricing on a direct Israel–Türkiye clash (directional anchor only).
- Contemporaneous reporting on the Syrian theater (the post-Assad order, the T4/Palmyra episode, and the Azerbaijan-hosted deconfliction channel).

Specific figures, dates, and attributions are retained in the companion evidence base and are available for inspection alongside this note.