METHODOLOGY NOTE

Levant Intervention Engine

Syria × Iranian Party in Lebanon (so-called Hizbullah) — model architecture, scenario logic, and calibration rationale

Author: Dr. Mosab Hawarey
URL: https://mosab.hawarey.org
Social: @DrHawarey (Twitter / LinkedIn)
Version: 1.1 · companion to the interactive simulator (single-file HTML, bilingual EN/AR)
Online: https://mosab.hawarey.org/levant-intervention-engine.html (English/Arabic toggle)
Date: 5 July 2026
Location: Amman, Jordan
Object modeled: the decision space opened by Washington's request that Damascus move militarily against Iranian Party (so-called Hizbullah) in Lebanon: one master decision (intervention depth), twelve drivers, five outcomes

AUTHOR'S NOTE — one scenario outside the model: the engine deliberately scores only unilateral Syrian action, because that is the question Washington has posed. The author's own assessment is that a structurally different configuration — a joint Syrian–Turkish–Jordanian force, acting on formal Lebanese requests to Damascus, Ankara, and Amman, with US and UK blessing — would sit outside this scenario space altogether: the coalition's combined capacity, legitimacy, and deterrent cover would remove the very constraints (no air power, rally-around-the-flag, third-party pile-on) that cap the unilateral case, making full demilitarization of the Iranian Party the expected outcome and positioning President al-Sharaa as a formidable regional figure. This is a reasoned judgment, not a model output; readers may weigh it accordingly.

This document explains how the engine 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 or armed actor.

Contents

1. Purpose and Scope

The engine maps the scenario space that opens if Syria's transitional government were to act militarily against Iranian Party in Lebanon, as Washington has pressed it to. It is the fourth tool in a family of geostrategic simulators — after the Islamabad MOU viability model, the Israel–Türkiye conflict model, and the Alawite feasibility index. Its distinguishing design is decision-conditionality: unlike the time-stepped Monte Carlo engines, this model conditions everything on one explicit choice — the depth of Syrian intervention, from none to a deep push into the Bekaa — and then reweights five qualitatively distinct outcomes as the analyst moves twelve drivers. It is deterministic: a given configuration always yields the same distribution, recomputed instantly.

Three properties were treated as design requirements: transparency (every outcome share is a traceable closed-form function of the depth choice and the named drivers), decision-legibility (the master decision is separated from the environment, so the analyst can see what each depth buys and costs under identical conditions), and falsifiability (every weight is stated explicitly so it can be disputed and re-tuned). The tool is a single bilingual file: an in-page toggle (and a ?lang=ar URL parameter) switches all text, layout direction, and numerals between English and Arabic while the engine and the analyst's current settings are unchanged.

What the tool is not: it is not a forecast, not an intelligence product, and not a recommendation for or against intervention. Its outputs are conditional on the analyst-set dials and on documented structural assumptions, and should be read directionally — as a disciplined way to reason about which factors move each outcome and by how much.

2. Model Architecture

Each configuration consists of a master decision $d \in \{0 \text{ None}, 1 \text{ Border}, 2 \text{ Shallow}, 3 \text{ Deep Bekaa}\}$ and twelve drivers on 0–100 scales, grouped as decision levers (US pressure, territorial inducement), Syrian-state factors (military capacity, internal cohesion), adversaries (Iranian Party's residual strength, Iran/IRGC reconstitution, Iraq/PMF hostility, Israeli aggression), and external backers (Türkiye, Gulf capital, Lebanese state/LAF strength).

Each of five outcome scenarios receives a raw score: score = depth-fit × scenario weight + condition term, where the depth-fit is a fixed four-entry vector expressing how compatible the scenario is with each intervention depth (Section 4), the scenario weight sets its structural prior, and the condition term is a signed weighted sum of the drivers relevant to that scenario (Section 5). The cascade and victory scenarios are additionally floored at zero inside the product and gated by depth; victory carries a structural cap (Section 6). All five raw scores are floored at 0.02 — no outcome is ever declared impossible — and normalized to sum to 100%. Five consequence metrics (casualties, regime stability, investment inflow, Iranian Party's degradation, regional-war risk) are then computed from the distribution and the drivers (Section 7).

3. The Five Outcome Scenarios

Scenario	Definition
A · Managed Neutrality	Damascus holds the border, interdicts smuggling, and backs the Lebanese state — no incursion. Tracks the revealed baseline strategy.
B · Limited Friction	Skirmishes and shallow cross-border sweeps against smuggling routes, then handoff to the LAF.
C · Deep Bekaa Quagmire	Syrian forces push into the Bekaa/Hermel and bog down under sustained fire — commitment without decision.

D · Catastrophic Cascade	Heavy losses, discredited government, decapitation risk, state fracture, regional spillover — the trap scenario.
E · Decisive Victory	Iranian Party broken, sanctions lifted, investment surges — the structurally capped best case.

The five span the qualitative space the open-source debate actually uses — from the "border regulator" reading of Damascus's current posture, through the quagmire warnings prominent in Arab commentary, to the low-probability transformation scenario invoked by intervention advocates. A/B are stability-side, C is the middle trap, D/E are the fat tails in either direction.

4. Depth-Fit Vectors and Gating

Each scenario's fit to the four depths (None, Border, Shallow, Deep) is fixed:

Scenario (weight)	None	Border	Shallow	Deep
A · Neutrality (×2.5)	1.00	0.90	0.35	0.10
B · Friction (×2.5)	0.15	0.45	1.00	0.50
C · Quagmire	0.00	0.10	0.45	1.00
D · Cascade	0.00	0.08	0.40	1.00
E · Victory (×0.8 cap)	0.00	0.05	0.35	1.00

The vectors encode the model's core structural claims. Neutrality is only fully available if the border is not crossed; friction peaks at shallow sweeps (a deep push converts friction into quagmire, hence B's fall from 1.00 to 0.50 at depth 3). Quagmire, cascade, and victory are all gated to zero at depth 0 — none of the deep-intervention outcomes, good or bad, is reachable without intervening — and open progressively with depth. Cascade admits a small opening (0.08) already at border depth, reflecting that even interdiction operations can be escalated from the other side. The casualty gate follows its own vector (0 / 0.15 / 0.50 / 1.00).

5. The Condition Terms

Each scenario's condition term is a signed weighted sum of the drivers (each scaled to 0–1). Signs and magnitudes, and why:

A · Neutrality: +0.7 Gulf +0.5 Türkiye +0.4 LAF +0.4 cohesion − 0.9 US pressure − 0.5 inducement. Neutrality is sustained by what Damascus stands to lose (Gulf capital, Turkish backing, a credible LAF handoff partner) and eroded by what pushes or tempts it across the border. US pressure carries the largest single coefficient in the stability-side terms because it is the named driver of the entire question.

B · Friction: +0.6 Syrian military +0.4 LAF − 0.5 Iran − 0.4 PMF. Contained friction requires enough capacity to conduct sweeps and a partner to hand off to; it is destabilized by the two actors most able to widen a skirmish.

C · Quagmire: +0.5 US pressure +0.4 inducement +0.3 Iranian Party's strength − 0.3 Syrian military. Quagmire is what happens when external push and temptation commit inadequate forces against a resilient defender — pressure and inducement raise it, capacity lowers it.

D · Cascade: +1.6 Israeli aggression +1.2 Iran +1.0 PMF +0.7 Iranian Party − 1.0 Syrian military − 0.9 cohesion − 0.6 Türkiye − 0.5 Gulf. The cascade term carries the model's largest coefficients because collapse is conjunctive and external: it is driven above all by third parties piling on — Israeli strikes on a committed Syrian force is the heaviest single accelerant — and braked by state capacity, internal cohesion, and the backers with money and skin in the game.

E · Victory: +1.5 Syrian military +1.0 cohesion +0.7 Türkiye +0.8 inducement − 1.3 Iranian Party − 1.1 Iran − 0.8 PMF − 0.6 Israel. Victory demands everything at once: a capable, united force, a backer, a motive — against a degraded enemy that no one resupplies and no third party complicates. The mirrored signs of D and E are deliberate: the same drivers that make collapse likely make victory unlikely, and the analyst can watch mass move between the two tails as they turn.

6. The Structural Victory Cap

The victory scenario is triple-restrained: its condition term sits on the smallest base (0.5, versus 1.0 for cascade and 1.5–2.5 for the others), the whole term is floored at zero, and the result is multiplied by 0.8. This encodes a structural judgment, not a sympathy: the new Syrian army has no air force and no armor mass, Iranian Party retains defensive depth in its own terrain, and a foreign invasion is the one event most likely to hand it a rally-around-the-flag recovery. Even with every dial at its best-case setting, victory therefore peaks in the high-thirties percent of the distribution rather than running away with it — the model refuses to let optimistic inputs manufacture a cheap decisive outcome, which is its sharpest single design statement.

7. Consequence Metrics

Five illustrative metrics are computed alongside the distribution:

- **Syrian troops killed:** casualty-gate(depth) × (0.4 + 0.5·Iranian Party + 0.4·Iran + 0.3·PMF + 0.4·Israel) × (1.3 − 0.6·Syrian military) × 3,000 — losses scale with depth, with the opposition actually met, and inversely with capacity. Zero at depth 0; up to ≈7,000+ at hostile extremes; rounded to the nearest 50 for display.
- **Regime stability (0–100):** 0.25 + 0.4·cohesion + 0.2·Türkiye + 0.15·Gulf − 0.5·(casualties/8000) − 0.4·P(cascade), clamped — stability is bought with cohesion and backing and spent in blood and cascade risk.
- **Investment inflow (0–$400B):** 0.05 + 0.45·Gulf·(stability/100) + 0.25·(P(A)+P(E)) − 0.4·P(cascade), scaled — Gulf capital is modeled as strictly stability-conditional, entering only when the board resolves toward neutrality or victory.
- **Iranian Party's degradation (0–100%):** 0.10 + 0.35·(P(C)+P(E)) + 0.25·(depth/3) + 0.2·LAF + 0.15·(1 − Iran) − 0.2·Iranian Party's strength — degradation accrues from contact, depth, a functioning LAF, and a cut resupply line; note that a quagmire degrades Iranian Party too, at a price the other metrics record.
- **Regional-war risk (0–100%):** 0.7·P(cascade) + 0.3·(depth/3)·(0.4·Iran + 0.4·PMF + 0.3·Israel) — dominated by the cascade share, with a depth-scaled term for third-party entry short of full cascade.

The metrics are illustrative magnitudes for reasoning about trade-offs — the casualty figure in particular is an order-of-magnitude device, not an estimate.

8. Input Variables, Defaults, and Presets

The Baseline 2026 preset encodes the open-source mid-2026 picture: Washington pressing, Damascus refusing, Iranian Party degraded but standing. "Direction" indicates whether raising the dial raises (▲) or lowers (▼) the cascade share at the baseline.

Variable	Default	Dir.	Basis (mid-2026 picture)
Intervention depth	1 · Border	▲	Interdiction/anti-smuggling — the posture Damascus has actually adopted.

US pressure intensity	65	▲	Sustained, public, and specific; the driver of the whole question.
Territorial inducement (Golan linkage)	20	▲	Floated in commentary; no concrete offer on record.
Syrian military capacity	35	▼	A new army without air cover or armor mass.
Internal cohesion	40	▼	Alawite, Druze, and Kurdish fractures unresolved.
Iranian Party's residual strength	35	▲	Heavily degraded 2024–26; command and rocket stocks reduced but real.
Iran / IRGC reconstitution	55	▲	Post-war reconstitution underway; resupply intent assumed.
Iraq / PMF hostility	55	▲	Militias positioned and rhetorically committed.
Israeli aggression vs. Syria	70	▲	High strike tempo and demonstrated decapitation posture.
Türkiye backing	65	▼	Defense relationship and stabilizing stake in Damascus.
Gulf investment flow	55	▼	Capital committed in principle, conditional on stability.
Lebanese state / LAF strength	45	▼	Disarmament process live but incomplete; partial handoff partner.

Two stress presets bound the space: Best case (deep intervention with every driver favorable) and Worst case (deep intervention with a weak, fractured force against fully mobilized adversaries). Presets are documented configurations for exploring the structure, not predictions.

9. Calibration and Anchoring

Structural anchor. The baseline is anchored to revealed behavior: Damascus has chosen border regulation over incursion, and the model's depth-1 default returns Managed Neutrality as the dominant outcome (≈65%), with Limited Friction (≈25%) absorbing most of the remainder and the deep-intervention outcomes (quagmire ≈5%, cascade ≈4%, victory <1%) held near their floors. That the tool, set to the real-world posture, reproduces the real-world reading — squeeze without invasion — is its primary calibration test.

Depth sweep. Holding the baseline environment fixed and moving only the master decision: at depth 0 neutrality is ≈90%; at depth 2 friction leads (≈44%) with quagmire and cascade at ≈16% each; at depth 3 cascade leads (≈37%) over quagmire (≈32%), with victory still ≈2% and expected casualties ≈4,000. The sweep expresses the model's central finding — under 2026 conditions, each additional increment of depth buys more downside than upside.

Preset bounds. The Worst case returns cascade ≈64% with regional-war risk ≈78% and stability at the floor; the Best case returns victory ≈38% — leading, but not dominant, against friction ≈26% and quagmire ≈25% — with the residual quagmire share and ≈1,300 expected deaths recording the cost even of success. These bounds are reported so the reader can see the full range the structure permits, and the asymmetry between them is the model's thesis stated numerically.

10. Validation

The engine was tested directly. The following checks all pass:

- The five outcome shares are non-negative and sum exactly to one across 300,000 randomized extreme configurations of all twelve drivers and all four depths — no out-of-range values, no NaN.
- The 0.02 floor guarantees every outcome remains reachable at every configuration; no scenario is ever displayed as impossible.
- All consequence metrics respect their bounds by construction (stability, degradation, and war risk in [0, 100]; investment in [0, 400]; casualties non-negative, ≈7,250 at the adversarial extreme).
- Depth gating behaves as designed: at depth 0, quagmire, cascade, and victory sit at their floors and casualties are exactly zero.
- Every driver moves the cascade share in its intended direction at the baseline — Israeli aggression, Iran, PMF, Iranian Party's strength, US pressure, and inducement raise it; Syrian capacity, cohesion, Türkiye, Gulf, and LAF strength lower it — verified by perturbation.
- Regional-war risk and expected casualties are monotone in depth at the baseline (0 → 10 → 24 → 45% and 0 → ≈600 → ≈2,000 → ≈4,050 respectively).
- The victory cap binds: under the Best-case preset, victory leads but does not exceed ≈38%, and the distribution is invariant under the language toggle — switching languages changes only presentation, never the numbers.

11. Limitations, Uncertainty, and Neutrality

Weights, not measurements. The depth-fit vectors, scenario weights, condition coefficients, and metric formulas 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, and the tool is designed for exactly that.

A deterministic engine. The model reweights scenarios; it does not simulate event sequences or time. Outcome shares are structural weights over qualitative futures, not frequencies from a stochastic process, and the consequence metrics are illustrative magnitudes.

Contested underlying data. Iranian Party's residual order of battle, Iranian reconstitution, Syrian force quality, and disarmament progress are all contested and partly opaque in open sources. Where they enter the model they do so as analyst-set relative magnitudes, not asserted facts.

Structure, not intelligence. The model encodes documented structure — the depth ladder, the actor set, the conditionality of Gulf capital, the third-party escalation routes. It has no access to classified information and infers nothing about any government's or armed actor's actual decision-making.

Neutrality. The engine privileges no outcome: the stability-side and escalation-side scenarios are scored by the same mechanism, cascade and victory are gated identically, and the assessment text is generated from the numbers rather than from a preferred narrative. No coefficient encodes a judgment that any party ought to act or refrain; the baseline's finding that non-intervention dominates is a property of the modeled structure under 2026 conditions, not a policy recommendation. This is a tool for thinking about a live and dangerous question, offered to clarify rather than to advocate.

Source Base

The model's inputs are drawn from the open-source record on the Syria–Lebanon file, mid-2025 to mid-2026. Principal categories:

- Contemporaneous reporting on US pressure on Damascus regarding Iranian Party, the Syrian government's public refusals, and the diplomatic sequencing around sanctions relief.
- Reporting and analysis on Iranian Party's condition after the 2024–2026 campaigns, the Lebanese disarmament track, and LAF capacity.
- Reporting on the new Syrian army's composition and limits, internal cohesion across communities, and the Turkish defense relationship.
- Reporting and commentary — notably in Arab media — on the quagmire/trap framing of a Syrian move into Lebanon, which the C and D scenarios formalize.
- Reporting on Gulf investment conditionality, Israeli strike posture toward Syria, Iranian reconstitution, and Iraqi militia positioning — the drivers of the condition terms.

Specific figures, dates, and attributions are retained alongside this note and are available for inspection. Where a magnitude is contested it enters the model only as a relative, analyst-set dial.