

METHODOLOGY NOTE

Alawite Breakaway State Feasibility Simulator

Model architecture, scoring logic, and calibration rationale

Author: Dr. Mosab Hawarey

URL: <https://mosab.hawarey.org>

Social: @DrHawarey (Twitter / LinkedIn)

Version: 1.0 · companion to the interactive simulator (single-file HTML, bilingual EN/AR)

Online: <https://mosab.hawarey.org/alawite-state-simulator.html>

Date: 5 July 2026

Location: Amman, Jordan

Object modeled: structural feasibility of an Alawite breakaway entity on the Syrian coast, scored today and projected over three horizons (5–10, 10–25, 25–50 years)

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, community, or movement.

Contents	Page
1. Purpose and Scope	3
2. Model Architecture	3
3. The Four Category Sub-Scores	3
4. The Blocking and Override Layer	5
5. Geographic Scope	5
6. The Classification Ladder	5
7. Time-Horizon Projection	6
8. The Precedent Radar	6
9. Input Variables, Defaults, and Presets	6
10. Calibration and Anchoring	7
11. Validation	8
12. Limitations, Uncertainty, and Neutrality	8
Source Base	9

1. Purpose and Scope

The simulator estimates the structural feasibility of a breakaway Alawite entity in coastal Syria — from cultural autonomy at the low end to a recognized independent state at the high end — as a function of twenty-nine analyst-set parameters. It is the third tool in a family of geostrategic simulators; the first modeled the viability of the Islamabad memorandum between Iran and the United States, and the second the Israel–Türkiye rivalry. Where those two are Monte Carlo path engines over a time trellis, this tool is a deterministic weighted-index model: every configuration of the dials maps to exactly one feasibility score, one scenario classification, and one set of horizon probabilities, recomputed instantly as the analyst moves a slider. There is no random draw and no seed; the model’s claim is structural, not probabilistic-path-based.

Three properties were treated as design requirements: transparency (the score is a traceable closed-form function of named inputs and published weights), comparability (the case is profiled against five secession precedents on identical axes), and falsifiability (every weight and default is stated explicitly so it can be disputed and re-tuned). The tool is fully bilingual (English/Arabic) with an identical engine in both languages.

What the tool is not: it is not a forecast, not an intelligence product, and not advocacy for or against secession. 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 feasibility and by how much.

2. Model Architecture

The engine has three layers. First, the twenty-nine inputs are aggregated into four category sub-scores — Internal, Legal/Normative, External/Geopolitical, and Trigger Conditions — each a weighted average on a 0–100 scale (Section 3). Second, the four sub-scores are combined into a base score using four meta-weights (default 20/10/40/30, auto-normalized; the analyst can re-weight them in an advanced panel). Third, and decisively, the base score is passed through a multiplicative blocking layer — parent-state consent, UNSC dynamics, anti-secession force proximity, geographic-scope difficulty, and the Türkiye/Hatay veto kill switch — each of which can only reduce the score (Section 4). The result is clamped to [0, 100].

The headline score is then translated into a scenario classification on an eight-rung ladder from Status Quo to Recognized Independent State (Section 6), and into three time-horizon probabilities (Section 7). A radar chart profiles the current configuration against five precedents (Section 8), and a sensitivity panel reports the five parameters whose +10 perturbation most moves the score (Section 11). Five preset scenarios let the analyst jump between documented configurations (Section 9).

The multiplicative design is the model’s central structural claim: secession feasibility is conjunctive, not additive. Strong internal cohesion cannot buy off a UNSC veto; generous external sponsorship cannot offset a parent state both unwilling and militarily present. This mirrors the empirical record — the cases that succeeded (South Sudan, Eritrea) cleared every gate at once, while cases strong on several axes but blocked on one (Somaliland) remain frozen.

3. The Four Category Sub-Scores

3.1 Internal factors (default meta-weight 20)

Component	Weight	Rationale
Alawite demographic cohesion	0.25	The community's territorial concentration and unity is the foundation of any claim.

Displacement volatility (inverted)	0.20	Sunni IDP inflow into Tartous dilutes the demographic base; the dial enters as (100 – value).
Economic viability	0.20	Port revenue, agriculture, and a functioning coastal economy.
Military capacity	0.15	Ability to defend the territory; also feeds the force-proximity modifier (\$4).
Governance capacity	0.10	Institutions able to administer a de facto entity.
Popular will / referendum potential	0.10	Demonstrable mass preference for separation, not merely elite ambition.

Cohesion and displacement carry the top weights because they are the slowest-moving and least substitutable ingredients: an entity without a coherent demographic core has nothing to secede with, however favorable everything else becomes.

3.2 Legal / normative factors (default meta-weight 10)

Component	Weight	Rationale
Remedial secession applicability	0.40	The strongest available doctrine: secession as remedy for grave, systematic persecution. The March 2025 coastal massacres are what give this argument any purchase at all.
Uti possidetis claim strength	0.30	Weak by design: no prior internal administrative boundary matches the claimed territory.
Referendum legitimacy	0.30	A credible, internationally observed vote is the normative price of admission.

The legal category carries the smallest default meta-weight (10) because international law is, in practice, the least binding constraint on secession outcomes: recognized cases succeeded politically first and were rationalized legally afterward. It is retained as a category because legal framing conditions how sponsors and blockers justify their stances.

3.3 External / geopolitical stances (default meta-weight 40)

Eleven state stances are set on a –100 (blocks) to +100 (sponsors) scale, mapped to 0–100 and combined with the following weights: Russia 0.25, United States 0.20, Türkiye 0.15, Israel 0.10, Iran 0.08, Saudi Arabia 0.05, France 0.05, China 0.04, United Kingdom 0.03, Germany 0.03, Jordan 0.02. Russia leads the weighting because of its naval and air presence on the coast itself (Tartus and Hmeimim) and its UNSC veto; the US is second as the system-level gatekeeper of recognition and sanctions relief; Türkiye is third as the neighbor with both the strongest motive and the greatest capacity to strangle a coastal entity. The external category carries the largest meta-weight (40) — the model's judgment, consistent with the precedent record, that recognition is granted from outside far more than it is earned from inside.

3.4 Trigger conditions (default meta-weight 30)

Component	Weight	Rationale
Syrian state collapse	0.30	The single event that most reshuffles feasibility: a functioning Damascus forecloses the question.
External sponsor commitment	0.30	A patron willing to underwrite the entity, not merely sympathize with it.

Regional war probability	0.15	War creates faits accomplis and distracted blockers — a risk-laden accelerant, not a stabilizer.
Great-power bargain	0.15	A Russia–US understanding that trades the coast into a wider settlement.
Alawite armed mobilization	0.10	Organized insurgency capacity; weighted last because mobilization without the other triggers invites suppression.

4. The Blocking and Override Layer

The base score is multiplied by four modifiers and one kill switch, each constructed so it can only hold or reduce feasibility, never raise it above the base:

- **Parent-state consent:** multiplier $0.15 + 0.85 \cdot (\text{consent}/100)$. At the default consent of 5, the multiplier is ≈ 0.19 — a non-consensual secession retains at most a fifth of its structural potential. Consent is the single heaviest lever in the model (Section 11), mirroring the bright line between South Sudan (consented, recognized) and Somaliland (unconsented, frozen).
- **UNSC dynamics:** multiplier $0.20 + 0.80 \cdot (\text{favorability}/100)$. Membership and collective recognition run through the Council; a hostile Council caps the ceiling near de facto status regardless of ground truth.
- **Force proximity:** multiplier $1 - 0.5 \cdot (\text{govForce}/100) \cdot (1 - \text{milCapacity}/100)$. Anti-secession forces on the doorstep suppress feasibility in proportion to how outmatched the entity is; the term vanishes as local military capacity approaches parity.
- **Geographic scope:** $\times 1.0$ (Core), $\times 0.7$ (PCCWS claim), $\times 0.3$ (Maximal) — see Section 5.
- **Türkiye/Hatay veto (kill switch):** $\times 0.02$ when the Maximal scope is selected and the veto toggle is on. Any claim touching Hatay converts Türkiye from opponent into existential adversary; the model treats this as effectively terminal and displays the switch as ACTIVE only in that scope.

The three continuous modifiers are floored above zero (0.15, 0.20, and 0.5 respectively) so that no single blocker mechanically zeroes the score — total foreclosure requires the kill switch or the compounding of several blockers, which is how the real cases behave.

5. Geographic Scope

Three mutually exclusive territorial definitions can be scored: Core (Latakia + Tartus governorates — the demographic heartland and both ports), the PCCWS claim (adding Homs and Hama, per the maximal coastal-council framing), and Maximal (adding northern Lebanon and Hatay). The scope multiplier (1.0 / 0.7 / 0.3) encodes that each expansion dilutes the Alawite demographic majority, lengthens indefensible frontiers, and multiplies the set of states with a direct stake in blocking the project. The Maximal scope additionally arms the Hatay kill switch. Scope is a selector, not a dial: the analyst is forced to commit to a territorial definition before reading the score, because feasibility claims that float across definitions are not falsifiable.

6. The Classification Ladder

The headline score maps to an eight-rung outcome ladder. The rungs are deliberately non-linear — narrow at the bottom, wide in the middle — because small early gains in feasibility translate into modest constitutional concessions long before they translate into statehood:

Score	Classification	Reading
< 5	Status Quo — No Change	The question does not seriously arise.

5 – 12	Cultural Autonomy Only	Language, religious, and local-administration concessions.
12 – 22	Administrative Decentralization	Devolved governorate powers within a unitary state.
22 – 38	Federal Region	Constitutional autonomy on the Iraqi-Kurdistan pattern.
38 – 55	De Facto Autonomous Entity	Self-governing in practice, unrecognized, contested.
55 – 70	Frozen Unrecognized State	Somaliland/Abkhazia pattern: durable, isolated.
70 – 85	Partially Recognized State	Kosovo pattern: substantial but incomplete recognition.
≥ 85	Recognized Independent State	South Sudan/Eritrea pattern: UN membership.

7. Time-Horizon Projection

The score is projected onto three horizons by fixed factors: 5–10 years = $0.35 \times \text{score}$, 10–25 years = $0.70 \times \text{score}$, 25–50 years = $1.15 \times \text{score}$, each capped at 95%. The monotone rise encodes the one strong regularity in the secession record: structural feasibility converts into outcomes slowly, and the constraints that dominate the near term (incumbent force posture, current alignments) decay faster than the enabling conditions (demography, geography, grievance). The long-horizon factor exceeds 1 because a structurally feasible project has more chances to catch a favorable shock over fifty years than over five; the 95% cap keeps every projection strictly short of certainty. These are horizon-conditional likelihood readings of a structural score, not survival curves from a stochastic process — the honest limit of a deterministic engine.

8. The Precedent Radar

The current configuration's five category profiles (Internal, Legal, External, Triggers, and Blocking Resistance — the averaged consent/UNSC/force modifiers expressed on 0–100) are plotted against five hard-coded precedent profiles: Kosovo, South Sudan, Eritrea, Somaliland, and Abkhazia. The precedent values are the author's structured judgments of each case at its decisive moment, on the same axes. The panel exists to discipline intuition in both directions: the successful cases show how high every axis had to be at once, and Somaliland — strong internally, near-zero externally — shows exactly which axis keeps a competent de facto state unrecognized for three decades. The comparison is qualitative and illustrative; no claim of case equivalence is made.

9. Input Variables, Defaults, and Presets

The default configuration (Preset 1, “Current Reality 2026”) encodes the open-source picture of mid-2026. Selected defaults and their basis:

Variable	Default	Basis (mid-2026 picture)
Demographic cohesion	45	Concentrated but shaken; post-March-2025 flight and emigration.
Displacement volatility	80	Heavy IDP inflow into Tartous is actively diluting the coastal balance.

Economic viability	30	Ports functioning; economy sanction-scarred and Damascus-dependent.
Military capacity	10	Ex-regime remnants disarmed or dispersed; no organized coastal force.
Governance capacity	15	No parallel institutions; local administration answers to Damascus.
Popular will	55	Real fear-driven separatist sentiment, short of a demonstrated mandate.
Uti possidetis	10	No prior administrative boundary supports the claimed unit.
Remedial secession	40	The 2025 coastal massacres give the doctrine live but contested purchase.
Referendum legitimacy	10	No mechanism, no observers, no prospect of a credible vote.
Syrian state collapse	20	The transitional state is consolidating, not collapsing.
External sponsor commitment	5	Sympathy exists; no state underwrites the project.
Regional war probability	25	Elevated background risk after the 2026 Iran war.
Great-power bargain	5	No coast-for-settlement trade is on any table.
Alawite armed mobilization	15	Sporadic insurgent activity, far from organized capability.
Parent-state consent	5	Damascus categorically rejects partition.
UNSC favorability	5	No permanent member advocates recognition.
Force proximity	85	Government-aligned forces deployed in and around the coast.

Country stances (–100 blocks ... +100 sponsors) default to: Russia –20, US –40, Israel 0, Türkiye –80, Iran –10, Saudi Arabia –50, France –30, UK –35, China –60, Jordan –45, Germany –40 — a board on which every weighty actor ranges from unhelpful to hostile, Türkiye and China anchoring the hard-block end (neighborhood stakes and anti-secession doctrine respectively), and Israel alone genuinely uncommitted.

Four alternative presets stress the structure: Syrian State Collapse (trigger-led), Russian Sponsor Pivot (patron-led), Israel Protection Umbrella (buffer-led), and Maximum Optimistic (every dial at its plausible best, including partial Damascus consent and the Hatay veto disarmed). Presets are documented configurations, not predictions.

10. Calibration and Anchoring

Structural anchor. The blocking layer makes near-zero the honest default: an unconsented secession, facing an unfavorable UNSC and proximate government forces, retains only a few percent of its base score by construction. This is the model expressing the international system's documented bias for territorial integrity, not a judgment about the community.

Precedent anchor. The five radar cases bracket the output space. The default configuration profiles far below every successful precedent on the External and Blocking axes — the two that the record says are decisive — and below Somaliland on Internal capacity.

Resulting default output. At the defaults, Core scope, the model returns a feasibility score of $\approx 0.7\%$ — Status Quo, with horizon probabilities of $\approx 0.3\%$ / 0.5% / 0.8% . The PCCWS scope returns $\approx 0.5\%$ and the Maximal scope $\approx 0.0\%$ (Hatay veto active). The stress presets return roughly 2–3% (Collapse, Sponsor Pivot, and Umbrella —

Cultural Autonomy territory at best) and $\approx 14\%$ under Maximum Optimistic (Administrative Decentralization, long-horizon $\approx 16\%$). The instrument's central finding is thus visible in its own bounds: even a conjunction of favorable shocks moves the outcome from “nothing” to “devolution,” not to statehood — full recognition requires parent consent and UNSC alignment that no plausible dial-setting of the 2026 board supplies.

11. Validation

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

- The feasibility score is a valid value in $[0, 100]$ across 200,000 randomized extreme dial configurations, all scopes and meta-weights — no out-of-range values, no NaN.
- Every parameter moves the score in its intended direction — verified by perturbation at the defaults: all sponsorship, capacity, legal, and trigger dials raise it; displacement volatility and government force proximity lower it; every country stance raises it as it moves from block toward sponsor.
- Meta-weights are auto-normalized: scaling all four by a common factor leaves the score unchanged.
- The blocking modifiers are individually bounded below (0.15, 0.20, 0.50) and can only reduce the base score; the Hatay kill switch binds only in the Maximal scope, exactly as labeled in the interface.
- Scope difficulty is monotone: for any fixed dial configuration, $\text{Core} \geq \text{PCCWS} \geq \text{Maximal}$.
- The classification ladder and horizon projections are total and consistent: every score maps to exactly one rung, and horizon values respect $\text{short} \leq \text{medium} \leq \text{long}$ with the 95% cap.

The sensitivity panel reports, for each parameter, the change in feasibility from a +10 perturbation holding all else fixed. Under the defaults the ranking is dominated by parent-state consent (+0.32 per +10) and UNSC favorability (+0.24), followed at an order of magnitude less by military capacity, force proximity, and the trigger dials — a hierarchy consistent with the model's central claim that the gates, not the aspirations, decide the question.

12. Limitations, Uncertainty, and Neutrality

Weights, not measurements. The component weights, blocking floors, scope multipliers, ladder thresholds, and horizon factors are reasoned judgments calibrated to the secession literature and the open-source record. They are the most contestable part of the model and are stated explicitly so they can be challenged; the advanced meta-weight panel exists so analysts can impose their own hierarchy.

A deterministic engine. The model has no stochastic layer: it scores structure and projects it, but does not simulate event paths. Horizon figures are conditional readings of the current structure, not survival probabilities, and carry no error bars of their own.

Contested underlying data. Coastal demographics, displacement flows, casualty figures from March 2025, and the internal politics of the community are contested and partly unknowable from 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 — demography, geography, doctrine, precedent, and the stance architecture of eleven states. It has no access to classified information and infers nothing about any actor's actual decision-making.

Neutrality. The tool takes no position on whether an Alawite entity should exist. Sponsorship and blocking are modeled symmetrically on the same scale; the grievance record enters only through the remedial-secession dial, where international law itself places it; and the default output — near-zero feasibility — is a statement about

the international system's architecture, not about the legitimacy of anyone's fears or aspirations. This is a tool for thinking about a question that is being asked in the region, offered to clarify rather than to advocate.

Source Base

The model's inputs are drawn from the open-source record on post-Assad Syria and from the comparative secession literature. Principal categories:

- Contemporaneous reporting on coastal Syria, 2024–2026: the fall of the Assad government, the March 2025 coastal violence and its documentation, displacement flows into Tartous, and the posture of the transitional authorities.
- Reporting and commentary on coastal political organization and claims, used to define the PCCWS scope option and the popular-will and mobilization defaults.
- The comparative secession and recognition literature — remedial secession, *uti possidetis*, and the recognition practice of the five radar precedents (Kosovo, South Sudan, Eritrea, Somaliland, Abkhazia).
- Open reporting on the stances of the eleven modeled states toward Syrian territorial integrity, and on the Russian coastal presence, Turkish red lines, and UNSC dynamics.

Specific figures, dates, and attributions are retained alongside this note and are available for inspection. Where a figure is publicly disputed — notably casualty and displacement counts — it enters the model only as a relative magnitude on an analyst-set dial.