METHODOLOGY NOTE

Islamabad MOU — 60-Day Viability Simulator

Model architecture, sequence logic, 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/islamabad-mou-simulator.html
Date: 19 June 2026 (v1) - 24 June 2026 (v2)
Location: Amman, Jordan
Window modeled: the 60-day MOU clock (signing → Day 60), six sequenced phases

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 models the viability of a specific, time-boxed agreement — the Islamabad Memorandum of Understanding — over its own 60-day negotiating window. 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 whether the Islamabad Memorandum of Understanding — the 14-point Iran–US framework signed on 17 June 2026 under Pakistani mediation to end the 2026 war — survives its own 60-day negotiating window, and with what quality of outcome if it does.[1] It is the first in a family of geostrategic Monte Carlo tools; a companion tool later modeled the Israel–Türkiye rivalry. The two share a visual grammar — Monte Carlo paths over a decision trellis — but this engine is a **viability assessment of a specific, time-boxed agreement** rather than a multi-year war forecast.

Three properties were treated as design requirements: **transparency** (every probability is a traceable function of named inputs), **structural fidelity** (the model walks the agreement's own sequence and paragraph structure rather than an abstract timeline), 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, and should be read directionally — as a disciplined way to reason about which factors move the deal's viability and by how much.

2. Model Architecture

Each simulated future is a path through six sequenced phases tied to the agreement's calendar: signing, the first compliance week, Day 15, the Day-30 ¶13 gate, Day 45, Day 50, and Day 60. The deal occupies one of three active **momentum bands** in any phase, with Return to War as an absorbing fourth state:

- **On Track** — implementation proceeding, the confidence-building measures landing. This is the start state at signing.
- **Strained** — slippage and friction; partial or delayed compliance.
- **Crisis** — broken measures, open recrimination, near-breakdown.
- **Return to War** — resumption of sustained hostilities; an absorbing terminal state.

In each phase the model first tests whether the deal collapses into war (Sections 3–4). If it does not, the band moves up or down according to a net momentum score (Section 6). At Day 60 the surviving paths resolve into one of three end-states — Durable Peace, Fragile Truce, or Cold Stall (Section 7). 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 a given dial configuration is reproducible, and the headline *Run* executes the full 10,000-path batch.

3. The Breakdown Mechanism: Six Sequential Phase Hazards

Unlike the companion model's noisy-OR over simultaneous pathways, breakdown here is **sequential**. At each phase a small Bernoulli "break" probability is drawn; if it fires, the path returns to

[1] Full released text of the 14-point memorandum (17 June 2026): CNN, Foreign Policy, and Wikisource. Paragraph numbering in this note follows that text.

war at that phase. The break probability is phase-specific and rises with the band — a deal already in Crisis is far likelier to rupture than one On Track. Two phases (P2, P5) carry an additional exogenous-spoiler test (Section 4). The sequencing mirrors the agreement's front-loaded, cliff-laden design: the dangerous moments are the early-compliance window, the Day-30 gate, the substantive negotiation, and the run-up to Day 60.

3.1 Phase summary

Phase	Window	What it tests	Break hazard (pre-band)	Primary drivers
P1 · Signing	Sign → D7	Do the immediate confidence-building measures begin?	0.010 (+0.04 in Crisis)	CBM compliance, mediator
P2 · Spoiler I	D7 → D15	A non-signatory shock breaks ¶1 from outside	spoiler sub-model (§4)	Israeli strike, Hezbollah, mediator, Trump follow-through
P3 · Day-30 gate (¶13)	D15 → D30	Have ¶¶1,4,5,10,11 begun? Do talks open?	0.018 (+0.06 Crisis) +0.04·snapback	CBM spine, Rubio, congressional snapback
P4 · Negotiation	D30 → D45	The deferred hard issues: enrichment, funds	0.020 (+0.09 Crisis / +0.015 else; −0.02 On Track) +0.04·IRGC +0.03·nuclear gap	Vance, Rubio, Pezeshkian, fund credibility, deadline; −nuclear gap, −IRGC, +Trump
P5 · Spoiler II	D45 → D50	A late non-signatory shock	spoiler sub-model + 0.03	as P2, elevated late in the clock
P6 · Outcome	D50 → D60	Final rupture, then horizon resolution	0.09 + 0.10·Israeli strike (Crisis only)	Israeli strike; then terminal resolver (§7)

Hazards are per-phase and pre-band; each is then scaled by the current band and clamped to a valid probability before the draw.

3.2 Why each phase is set as it is

P1 — Signing (lowest base, 0.010). The ceasefire held at signing; the near-term risk is non-compliance, not rupture, so the unconditional break is small. A path already in Crisis carries an added 0.04, reflecting that even the opening week is combustible once trust has broken.

P2 and P5 — Spoiler windows. ¶1 names Lebanon explicitly precisely because Israeli action there nearly collapsed the April ceasefire; Israel and Hezbollah are non-signatories who can shatter ¶1 from outside without either party defecting.

[2] Two windows are modeled because spoiler risk is continuous, and P5 carries an extra +0.03 as patience thins late in the clock.

P3 — The Day-30 ¶13 gate. This is the structural hinge of the agreement. ¶13 makes the opening of substantive talks conditional on the beginning of ¶¶1, 4, 5, 10 and 11.

[2]¶1 names Lebanon explicitly because Israeli action in Lebanon repeatedly strained the April 2026 ceasefire; Israel and Hezbollah are non-signatories able to break the front from outside (Wikipedia: 2026 Iran war ceasefire; 2025–2026 negotiations).

[3] The model honors this: if the compliance spine has not begun (the CBM index below 0.45), momentum takes a hard penalty at the gate, and snapback risk both raises the break and drags momentum down.

P4 — Negotiation. Where the deferred, hardest issues live: enrichment (¶8, "agree to agree")[4] and the credibility of reconstruction (¶6), under the 60-day deadline (¶3). IRGC resistance and the nuclear gap raise the break; the dealmaker coalition (Vance, Rubio, Pezeshkian) and a credible reconstruction package lower it. The On-Track band receives a −0.02 discount — a deal that reaches negotiation healthy is structurally safer.

P6 — Outcome. A final Crisis-band rupture driven by Israeli action is tested; otherwise the path resolves into a terminal end-state (Section 7).

4. The Exogenous-Spoiler Sub-Model

The spoiler hazard is the model's sharpest structural point: the signatories can comply in good faith and still be dragged back to war by actors outside the agreement. Its raw intensity is **0.6·(Israeli strike) + 0.4·(Hezbollah)** — Israel weighted higher as the more capable and more likely initiator. That raw figure is scaled by the band and by the "spoilers alone can trigger war" toggle:

- **Toggle on (default):** On Track ×0.22, Strained ×0.52, Crisis ×0.90 — a spoiler can break the deal even from a healthy state, but the odds climb steeply as the relationship deteriorates.
- **Toggle off:** On Track ×0.10, Strained ×0.22, Crisis ×0.50 — spoilers bite meaningfully only once the deal is already wobbling.

The result is then dampened by Pakistani mediation — ×(1 − 0.30·mediator) — and by Trump follow-through — ×(1 − 0.18·Trump) — with P5 adding a small fixed increment. Mediation and sustained presidential ownership are thus the two levers that most directly suppress the route to war that neither signatory controls.

5. The ¶13 Compliance Spine and the Day-30 Gate

Four compliance dials — the US lifting the naval blockade (¶4), Iran demining and guaranteeing Hormuz passage (¶5), US oil-export waivers (¶10), and the release of frozen funds (¶11) — average into a single confidence-building (CBM) index. ¶13 conditions the opening of substantive talks on these measures (together with ¶1) having begun and continuing.

The model treats compliance, not negotiation, as the early center of gravity. The CBM index drives momentum at P1 and P3, and a low index (below 0.45) imposes a momentum penalty at the Day-30 gate. The spine is front-loaded and asymmetric: the US acts first and most visibly — blockade, waivers, funds — while Iran's demining is the reciprocal up-front lever, and Hormuz free passage is written as "60 days only," a cliff rather than a settlement.

---

[3] ¶13 conditions the opening of substantive talks on the beginning and continued implementation of ¶¶1, 4, 5, 10 and 11 — the confidence-building spine modeled here (CNN; Iran International text).
[4] The text defers the nuclear question and refers to a "minimum methodology" for neutralizing the highly-enriched-uranium stockpile via in-country down-blending after a final deal — the "agree-to-agree" gap (CNN).

6. Band Dynamics and Momentum

When the deal does not break in a phase, the band moves according to a net momentum score: a phase-specific weighted sum of the stabilizing factors (the compliance spine, mediation, the dealmaker coalition, fund credibility, deadline discipline) minus the escalatory ones (the nuclear gap, IRGC resistance, congressional snapback, and spoiler pressure). A positive score makes a move toward On Track more likely; a negative score makes a move toward Crisis more likely; the larger the magnitude, the stronger the pull.

The functional form — a base 0.34 probability of improving or worsening, shifted by 0.42×momentum — is intentionally simple and saturating: extreme dial settings push the move probabilities toward certainty without ever exceeding it. This is the form later inherited by the companion model. Different phases read different drivers: the early and spoiler phases lean on compliance and mediation; the Day-30 gate on the spine, Rubio, and snapback; the negotiation phase on the full deal vector.

7. Terminal Resolution at Day 60

At Day 60, each surviving path resolves into Durable Peace, Fragile Truce, or Cold Stall. The split depends on the band the path occupies and on a quality score drawn from the structural dials: $q = -1.6 \cdot (\text{nuclear gap}) + 0.8 \cdot (\text{fund credibility}) + 0.8 \cdot (\text{deadline discipline}) + 0.4 \cdot (\text{mediator})$, each centered at its mid-point. An On-Track path resolves overwhelmingly to Durable Peace; a Strained path mostly to Fragile Truce, with tails either way; a Crisis path mostly to Cold Stall. The three probabilities are constructed to be non-negative and to sum to one for every band and every dial setting — a property verified directly (Section 10).

This resolver produces the reported terminal probabilities in the Monte Carlo run. In the interactive *walk-the-tree* mode, where the analyst drives a single path by hand, the Day-60 node is assigned by the band the path has reached; all four end-states, including Return to War, remain reachable in that mode as well.

8. Input Variables and Default Settings

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

Variable	Default	Dir.	¶	Basis (late-June 2026 picture)
US lifts naval blockade	0.68	▼	4	Removal within 30 days; the first visible US obligation.
Iran demines / Hormuz passage	0.62	▼	5	Safe commercial transit — reciprocal up-front lever; “60 days only.”
US oil-export waivers	0.60	▼	10	Treasury licenses crude & banking; waived, not eliminated.
Frozen-funds release	0.52	▼	11	Assets to the Central Bank of Iran; procedures still to be agreed.
Israeli strike likelihood	0.28	▲	1	Non-signatory; can break the Lebanon front from outside.

Variable	Default	Dir.	¶	Basis (late-June 2026 picture)
Hezbollah / Lebanon flare-up	0.22	▲	1	Re-ignites a named front without either party defecting.
Vance posture (hawk→dealmaker)	0.55	▼	—	Signed for the US; sustained ownership cuts both ways.
Rubio backs the deal	0.33	▼	—	Historic Iran hawk; may slow-walk waivers.
Trump follow-through	0.55	▼	1	Sustained presidential attention dampens spoilers.
Congressional snapback risk	0.40	▲	7	New sanctions quietly reverse the relief.
IRGC / hardliner resistance	0.45	▲	8	Pushback on demining and the nuclear freeze.
Pezeshkian govt durability	0.58	▼	2	Capacity to deliver Iran's commitments.
Nuclear "agree-to-agree" gap	0.62	▲	8	Enrichment redline vs. on-site down-blending.
$300B credibility	0.45	▼	6	Reconstruction pledge believable — the figure is publicly contested.
60-day deadline discipline	0.60	▼	3	Momentum vs. drift on a hard, extendable clock.
Pakistan mediator leverage	0.55	▼	12	Monitoring and good offices; Qatar also facilitating.

The $300B reconstruction figure is modeled as a credibility dial rather than a fixed sum, because the number itself is publicly contested.[5]

9. Calibration and Anchoring

Because the phase hazards and weights are priors rather than measured frequencies, the aggregate output is anchored to the open-source record in two ways.

Structural anchor. The agreement's own design makes "signed but fragile" the honest default: obligations are front-loaded onto the US, nothing is binding until a final deal endorsed by a UNSC resolution (¶14), Hormuz free passage is a 60-day cliff (¶5), and the hardest issue — enrichment — is explicitly deferred (¶8). A structure this conditional cannot calibrate to a high-confidence peace.

Consensus anchor. Contemporaneous reporting supports a default with substantial mass on the sub-war outcomes alongside a real, double-digit return-to-war tail: G7 governments welcomed the framework, while analysts described it as fragile, Iranian hardliners and the IRGC pushed back hard against the negotiators,[6] US congressional and intra-coalition hawks raised snapback pressure,[7] and the reconstruction figure was openly disputed.

Resulting default output. At the default settings, with spoilers able to trigger war, the model returns approximately: Durable Peace 25%, Fragile Truce 30%, Cold Stall 17%, and Return to War 28% — a peace mass (~55%) modestly ahead of breakdown. Disabling lethal spoilers lowers war to roughly 22%. Pushing all dials to their safest settings drives war to about 7% (an irreducible early-break

[5]The roughly $300B reconstruction figure is contested in public reporting; it was reported as part of the framework and subsequently disputed by President Trump, which is why credibility — not the headline number — is the modeled dial (Wikipedia: 2025–2026 negotiations).
[6]Iranian hardline and IRGC-aligned backlash against the negotiators is documented in contemporaneous reporting (Iran International).
[7]US congressional and intra-coalition ("MAGA hawk") opposition and snapback pressure are documented across contemporaneous reporting (Al Jazeera, Axios, The Hill, The New Republic).

residual); pushing them to their most dangerous drives it above 97%. These bounds are reported so the reader can see the full range the structure permits.

10. Validation

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

- Every phase break probability and spoiler hazard is a valid probability in [0,1] by construction (each is clamped before the draw); no out-of-range values or NaN arise under randomized extreme dial settings.
- Only valid end-states are ever produced; terminal counts sum exactly to the number of simulated paths — no lost or double-counted futures.
- Return-to-war transitions reconstructed from the path edges sum exactly to the War total — internal accounting is consistent.
- Return-to-war risk is monotonic in the compliance spine: it never rises as any of the four ¶13 dials (¶¶4, 5, 10, 11) is increased — verified across full sweeps of each.
- The break hazard increases strictly with the band (On Track < Strained < Crisis) by construction — every band-dependent term is additive and non-negative.
- All sixteen dials move return-to-war risk in their intended direction — verified exhaustively; the five risk dials raise it, the eleven stabilizers lower it.
- The terminal resolver's three outcome probabilities are non-negative and sum to one for every band and dial setting — confirmed over 200,000 randomized configurations.
- All four outcomes — including Return to War — are reachable in both the Monte Carlo run and the interactive walk-the-tree mode.

The sensitivity panel reports, for each variable, the change in cumulative return-to-war probability from a one-notch move toward danger, holding all else fixed. Under the defaults the largest movers are the two spoiler dials (Israeli strike and Hezbollah), followed by the negotiation drivers — Rubio, IRGC resistance, the nuclear gap, and Vance — consistent with a model in which the route neither signatory controls is the dominant danger.

11. Limitations, Uncertainty, and Neutrality

Priors, not measurements. The phase hazards, spoiler multipliers, and momentum 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, and the tool is designed for exactly that.

A short, sharp horizon. The 60-day window is the agreement's own, and the model does not reach past it. Its claim is about whether the framework survives to a final deal, not about what any final deal would contain or how it would hold.

Contested underlying data. Compliance status, the reconstruction figure, and the spoiler tempo all vary by source and are sometimes partisan. Where they enter the model they are treated as relative magnitudes and as dials, not precise values.

Structure, not intelligence. The model encodes documented structure — the paragraph sequence, the cliffs, the deferred issues, and the non-signatory spoilers. It has no access to classified information and infers nothing about either side's actual decision-making.

Neutrality. The agreement, the US, and Iran are each given their own SWOT lens on equal terms. No dial or default encodes a judgment that one party is acting in bad faith. This is a tool for thinking about an agreement's viability, offered to clarify rather than to advocate.

Source Base

The model's inputs are drawn from the released text of the Islamabad Memorandum and contemporaneous open-source reporting from June 2026. Principal categories:

- The agreement text: the 14-point memorandum as released on 17 June 2026 (CNN, Foreign Policy, Wikisource), which fixes the paragraph structure the model walks.
- Contemporaneous reporting on signing, mediation, and the post-signing period (Al Jazeera, Iran International, Reuters, Wikipedia compilations).
- Reporting on the domestic backlash on both sides — Iranian hardline/IRGC opposition and US congressional and intra-coalition resistance — used to set the relevant dials.
- Reporting on the Lebanon spoiler dynamic and the April-2026 ceasefire strains, used to justify the exogenous-spoiler sub-model.

Specific figures, dates, and attributions are retained alongside this note and are available for inspection. Where a figure is publicly disputed — notably the reconstruction sum — it is modeled as a credibility dial rather than asserted as fact.