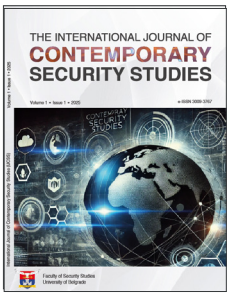
Volume 2 • Issue 2 • 2026

Faculty of Security Studies, University of Belgrade

International Journal of Contemporary Security Studies (IJCSS)



Article

The Credibility Cliff: AI-Enabled Anti-Submarine Warfare and Anticipatory Posture Change in Asymmetric Nuclear Dyads

Aleena Saif Ullah[1,*]

[1] University of the Punjab, Department of Political Science (Quaid-i-Azam Campus, Lahore 54590), Pakistan; aleenasaifullah68@gmail.com.

[*] Correspondence: aleenasaifullah68@gmail.com; tel.: +92-301-3074454.

Received: 11 May 2026; Revised: 10 September 2026; Accepted: 1 October 2026; Published: 5 November 2026.

ABSTRACT

How do states behave when an emerging counterforce technology begins to erode the survivability of their nuclear deterrent before that erosion can be verified? This article develops two formal instruments for answering that question and illustrates them through AI-enabled anti-submarine warfare in the China–India–Pakistan strategic triangle. The credibility cliff is derived as the aggregated neutralization probability at which a state's modeled deterrence condition ceases to hold, with retaliatory capability represented additively across sea-based and residual land and air legs. Its dual, the residual retaliation requirement, measures the survivability burden that adversary capability transfers onto the residual legs, and rises at a rate equal to the sea leg's undegraded deliverable capacity. The Posture Reversion Mechanism derives, from expected-utility maximization over bounded beliefs, a reversion threshold preceding the structural threshold by an anticipation margin that increases with belief uncertainty. Parameters are calibrated from categorized open sources with an evidence cut-off of 31 December 2025, and Monte Carlo simulations propagate all author-assigned values alongside complete deterministic grids. India's China-mission sea leg is cliff-exposed in 54.2 percent of draws (median threshold 0.45) and structurally insufficient in a further fifth; China is cliff-immune in every draw; and Pakistan, contrary to an earlier formulation of this framework, is cliff-immune under the sourced central survivability scenario, its sea-based program functioning as a hedge against declining land-leg survivability rather than as the foundation of assured retaliation. A two-player China–India interaction with Pakistan as the externality bearer yields, conditional on one explicitly named preference assumption whose role is quantified by the complete ordinal enumeration, an equilibrium that bilateral agreement does not resolve. The central finding is that AI-enabled anti-submarine warfare does not abolish assured retaliation but shifts the survivability burden at a calculable rate and devalues sea-based hedges before it disables sea-based deterrents.

KEYWORDS

Nuclear deterrence; anti-submarine warfare; second-strike capability; anticipatory posture change; artificial intelligence; asymmetric nuclear dyads; strategic stability.

ONLINE FIRST

1. Introduction

How do states behave when an emerging counterforce technology begins to erode the survivability of their nuclear deterrent before that erosion can be empirically verified? This question sits at the intersection of three literatures that have not been fully joined: the classical work on second-strike

Academic Editor:
Prof. Dr. Vladimir M. Cvetković
Copyright: © 2026 by the authors.

Ullah, A. S. . (2026). The Credibility Cliff: AI-Enabled Anti-Submarine Warfare and Anticipatory Posture Change in Asymmetric Nuclear Dyads.*International Journal of Contemporary Security Studies*, 2(2), 99-120.

survivability, the contemporary analysis of counterforce modernization, and emerging assessment of artificial intelligence in military sensing.

The problem is consequential because capability degrades continuously while posture changes discretely, and posture changes in anticipation. For a state whose retaliatory capability depends substantially on the single sea-based delivery leg, or which maintains a sea-based hedge against the erosion elsewhere, the interval preceding any threshold-crossing may be more dangerous than crossing itself, since one state may already have shifted toward the higher escalation risk while its adversary still believes deterrence intact.

In August 2025, a research team at China Helicopter Research and Development Institute published an account of an AI-enabled anti-submarine warfare decision system reporting a reduction in the simulated submarine survival from eighty-five to five per cent (Wang et al., 2025), reported subsequently in the international coverage (South China Morning Post, 2025). These are simulation outputs under the modeled conditions. They do not demonstrate operational performance against a real submarine in the contested environment, and they enter the present analysis only as an upper-bound input. What can be characterized more confidently is the architectural direction the work represents: simultaneous fusion of acoustic, radar, thermal, magnetic-anomaly, and oceanographic data at machine speed, treated as the design principle rather than as an augmentation of acoustic detection. The same architectural direction is emphasized independently in the North Atlantic Treaty Organization assessments of artificial intelligence in the anti-submarine warfare (Michell, 2021), is visible in Western maritime-defense mission systems (Ultra Maritime, 2025), and is discussed in the recent technical commentary on autonomous maritime sensing (Stupples, 2025).

1.1. Prior Literature and the Gap Addressed

Existing frameworks identify the conditions under which the second-strike survivability is threatened. Wohlstetter (1959) established that survivability is a demanding and contingent achievement rather than just considering it an automatic consequence of possession. Lieber and Press (2017) demonstrate that improvements in the accuracy and the remote sensing have shifted strategic balance toward counterforce feasibility. Long and Green (2015) show that intelligence capability, rather than the warhead numbers alone, drives counterforce prospects. Powell (2003) formalises interaction between the defensive capability and the deterrence stability, and Narang (2014) establishes that the regional nuclear powers adopt structurally distinct postures shaped by capability and threat environment. Sagan (1994) and Mueller (1988) supply organisational and sceptical counterweights to capability-driven accounts. On specific question of artificial intelligence and counterforce, Loss (2019) argues that machine processing of sensor data may supply missing element in the counterforce chain, this is an argument that the present analysis formalises and bounds.

The claim advanced here is deliberately narrow. This is not first recognition that states adapt in anticipation of capability shifts; that dynamic is present in Wohlstetter's original formulation, is central to Long and Green's account of the counterforce intelligence, and is implicit in survivability literature generally. The contribution is a particular formalisation: anticipatory posture change expressed as a threshold condition with calculable comparative statics, a decision-theoretic derivation of the anticipation margin from stated utilities, a dual instrument rendering the transferred survivability burden measurable, and a transparent calibration in which every contestable parameter is declared and varied.

Recent scholarship supports two premises of this framing. Analysis of deterrence against a nuclear-armed adversary in the Ukraine war underlines the limits of classical bilateral frameworks under conditions of asymmetry (Edmond, 2026). Work on multi-actor deterrence demonstrates that the deterrence problems of small and mid-sized states are rarely reducible to a single dyad (Pavić & Beriša, 2026; Cvetković, 2026), which is precisely the structural situation formalized in Section 3.6.

1.2. Objectives and Propositions

The research question is: at what point in the development of an adversary's AI-enabled anti-submarine warfare capability does a state with a predominantly sea-based deterrent, or a sea-based hedge, rationally alter its nuclear posture, and what observable behavior does that alteration produce?

Three falsifiable propositions follow.

1. P1 (Anticipation). Posture change, coded through the indicator framework in Section 3.7, occurs before any publicly verified capability demonstration crosses the relevant structural threshold. This is disconfirmed if indicators move only after verified demonstrations, or move independently of anti-submarine warfare developments.
2. P2 (Asymmetry). Holding sea-based force size and delivery probability constant, states with smaller residual retaliation capability reach their reversion threshold at lower perceived neutralization probabilities. This is disconfirmed if reversion timing is uncorrelated with residual capability.
3. P3 (Corner solutions and hedges). A state whose deterrence condition fails at zero neutralization probability exhibits no reversion dynamics but a stable corner solution in which asymmetric-escalation doctrine persists independently of adversary capability milestones, while a cliff-immune state maintaining a sea-based hedge exhibits hedge-foreclosure behavior as capability approaches its hedge threshold. This is disconfirmed if a structurally insufficient state sustains the assured retaliation, or if a hedging state's investment is insensitive to the perceived capability trends.

1.3. Scope Conditions

The framework applies where four conditions hold jointly: the state's assured-retaliation aspiration rests primarily on a sea-based leg, or it maintains a sea-based hedge against declining survivability elsewhere; an adversary is investing visibly enough in AI-enabled ASW to generate public indicators on which beliefs can update; the survivable, deliverable contribution of residual legs sits below or near the deterrence floor, making sea-leg attrition strategically meaningful; and leadership acts on open indicators rather than awaiting demonstrated capability.

Where conditions fail, predictions attenuate. Diversified great powers fail the third, which is why China is cliff-immune throughout. States under credible extended-deterrence guarantees substitute allied capability for their own and fall outside the first; the alliance-coordination dynamics governing such cases (Mulyanyuma, 2026) lie beyond this framework. One qualification matters especially: the framework addresses a mission-specific second-strike component, not a state's retaliatory capability as a whole. India is assessed only on its China-directed mission; its Pakistan-facing deterrent does not depend on the sea leg.

One caveat bounds the technology mechanism throughout. The trajectory from published simulation to fielded capability is neither linear nor inevitable. Counter-detection, platform quieting, decoy proliferation, cyber and electronic attack on sensor networks, adversarial manipulation of training data, and unresolved engineering limits of multi-domain fusion under wartime conditions can each flatten or reverse capability curve. The argument does not require technology to succeed; it requires only that target states believe it might, and update posture on that belief.

2. Methods

2.1. Study Design and Rationale

The study is formal-analytical, bringing together closed-form algebraic derivation, decision-theoretic modeling, probabilistic sensitivity analysis, and structured case comparison. What is at issue is a threshold relationship between adversary capability and posture choice; one that resists statistical

estimation because the dependent variable is rare, partly unobservable decision, but that can be derived formally with its parameter dependence laid open completely. The governing methodological commitment is to declare and propagate uncertainty rather than produce point estimates.

2.2. Setting and Data Sources

The empirical setting is the China–India–Pakistan nuclear triangle from 2016 to 2025, with an evidence cut-off of 31 December 2025. All sources are unclassified and publicly accessible: peer-reviewed scholarship, the Nuclear Notebook series (Kristensen, Korda, & Johns, 2023; Kristensen, Korda, Johns, & Knight, 2024), the Pentagon's annual report on Chinese military developments (U.S. Department of Defense, 2024), the IISS (2024) Military Balance, official government press releases (Inter-Services Public Relations, 2017), established defence reporting, and GEOnet Names Server (National Geospatial-Intelligence Agency, n.d.) for geographic coordinates.

2.3. Data Collection and Study Units

Three states that possess or are developing sea-based nuclear delivery capability serve as study units; they are selected on theoretical grounds rather than sampled. They, together form single interacting strategic system, whose structural parameters are estimable from unclassified material, and they vary along the dimensions the framework treats as decisive: residual retaliatory depth, propulsion type, and geographic constraint, permitting differentiated rather than confirmatory application of the model.

We classified every empirical input into one of seven evidence categories before use: confirmed, officially reported, independently assessed, estimated, projected, simulated, and assumed. The parameter tables record the category alongside the value and source. The practice of describing all open-source inputs as verified, adopted in an earlier version of this analysis, has been discontinued.

Faculty of Security Studies
University of Belgrade

2.4. Ethics and Compliance

The study involved no human participants, no animal subjects, and no personal data, and therefore required no ethical approval. All source material is publicly available and unclassified; no classified information was accessed or used at any point. Section 4.4 reports a dual-use assessment governing the aggregation and presentation of open-source military information.

2.5. Data Processing and Analysis

Analysis proceeded in four steps. First, the deterrence condition and its two dual instruments were derived in closed form, with each parameter defined, decomposed into components, and assigned an evidence category before use. Second, the Posture Reversion Mechanism was derived from expected-utility maximisation over beliefs held on the bounded support from zero to one, using a beta belief family, yielding a probability threshold and a closed-form anticipation margin. Third, all author-assigned parameters were propagated by Monte Carlo simulation using 200,000 draws per configuration with uniform independent sampling across full declared ranges, cross-checked against complete deterministic grids; neutralisation was simulated at platform rather than warhead level using a one-factor Gaussian copula so that correlated exposure is represented explicitly. Fourth, the strategic interaction was analysed as a two-player normal-form game with a third-party externality, and its equilibrium was tested by complete enumeration of all 576 strict ordinal preference profiles rather than by comparison against selected alternative matrices.

Analyses were implemented in Python 3.11 using NumPy 1.26 and SciPy 1.11. The pseudo-random number generator was initialised with seed 20260722 for all reported results. Complete code, equations, coordinates, and grids are provided in Appendix A, sufficient for full replication.

3. Results

3.1. The Technology Mechanism and the Snorkel Exposure Index

The architectural significance of current development lies in the treatment of multi-domain fusion as the core design principle. Acoustic detection has historically been the limiting modality; fusion architectures instead combine weak signals across modalities, so that countermeasures optimised against any single channel provide diminishing protection. Reported adaptive behaviours, including responses to zigzag evasion and decoy deployment, describe search-policy optimisation rather than any single-sensor breakthrough.

The Snorkel Exposure Index quantifies the fraction of patrol time during which a conventionally powered submarine must operate at or near the surface for air induction:

$$\mathrm{SEI} = (f_s \times d_s) / 24$$

At one snorkelling event per day and four hours per event, the index equals 0.167. Three clarifications govern its use. First, the index is an exposure-time measure and not a detection probability. It bounds the opportunity available to detection modalities that depend on snorkelling, including radar, visual, infrared plume, and near-surface acoustic signature, the last of which is increasingly complemented by advances in magnetic-anomaly sensing (Yu et al., 2025), and enters the model as one contributor to the aggregated neutralisation probability rather than as that probability itself. No functional mapping from the index to detection probability is asserted, because none is available at the unclassified level. Second, nuclear-powered submarines have zero snorkel exposure but not zero detectability; they do not snorkel, yet remain detectable through radiated acoustic noise, magnetic anomaly, thermal wake, hydrodynamic surface signature, and space-based observation. Third, the snorkelling parameters derive from Cold War-era literature on diesel-electric operations (Stefanick, 1987; Friedman, 2006). Air-independent propulsion substantially extends submerged endurance, so one event per day is best read as a conservative upper bound for such a boat on extended patrol, and the resulting index as an upper bound on snorkel-channel exposure. This is declared as an assumption rather than defended as established.

3.2. The Deterrence Condition

Let p denote the adversary's aggregated per-patrol probability of successful neutralisation of a sea-based warhead's retaliatory contribution, bounded between zero and one. This is deliberately not a detection probability. The counterforce chain runs from detection through classification, continuous tracking, targeting, engagement, destruction, battle-damage assessment, command-and-control disruption, and prevention of launch before completion; failure at any stage, including launch after detection, breaks the chain. The parameter p is the compound probability that the entire sequence succeeds, and is bounded above by the probability of any single stage. Published simulation detection rates therefore enter as upper-bound inputs, not as estimates of p.

Let N_{sea} denote nuclear warheads actually deployed on operational patrol aboard sea-based platforms. This is not the number of submarines, missile tubes, or warheads notionally assigned to the naval force; these quantities are not interchangeable, and Tables 1 and 2 state which quantity each source supports.

Let $P(\mathrm{D})_{sea}$ denote the probability that a surviving sea-based warhead is successfully delivered, decomposed as the product of reach, reliability, and penetration. Reach is a binary geographic condition computed from missile range against haversine target distances; reliability is the probability

of successful launch and flight; penetration is the probability of surviving adversary defences. Only reach derives from geography. Range establishes whether a target is attainable, but does not establish reliability, launch success, defence evasion, or arrival, and the decomposition separates a necessary condition from a sufficient one.

Let A denote the residual retaliation term, the expected number of surviving and deliverable non-sea warheads available after an adversary counterforce first strike, formed as the sum over land and air legs of the product of warhead count, leg survival probability, and leg delivery probability. This term represents land- and air-based components as separate, additive contributors. It replaces the multiplicative triad-depth factor used in an earlier version of this framework, and the change is substantive rather than notational: under a multiplicative specification the total is the product of sea-based capability and a triad-depth score, so that a low score does not merely fail to credit the land leg but reduces the sea leg's contribution. For a state whose arsenal is predominantly land-based this inverts the intended representation. The additive form eliminates the inversion, and because N_{sea} and $P(\mathrm{D})_{sea}$ enter exactly once while A covers only non-sea legs, it eliminates by construction the double-counting to which a multiplier is vulnerable.

Expected retaliatory capability is therefore

$$\mathrm{E}\underline{R} = N_{sea} \times (1 - p) \times P(\mathrm{D})_{sea} + A$$

and, with D_{min} the minimum deterrence floor, the deterrence condition is that expected retaliatory capability is at least the floor.

Solving at equality yields the credibility cliff

$$p^{\star} = 1 - (D_{min} - A) / (N_{sea} \times P(\mathrm{D})_{sea})$$

and solving instead for the residual capability required at a given neutralisation probability yields the residual retaliation requirement

$$A^{\star}(p) = \max\{0, D_{min} - N_{sea} \times (1 - p) \times P(\mathrm{D})_{sea}\}$$

These are duals of the same condition: evaluating the requirement at the cliff returns A exactly. The requirement is treated as the principal instrument for three reasons. It is defined across the whole probability interval, whereas the cliff is undefined or uninformative in two of the three cases below. It expresses consequence of adversary capability growth in units a force planner can use, and pins down what the technology does to nuclear posture. AI-enabled ASW does not abolish retaliation; it shifts the survivability burden onto whatever legs remain at a fixed rate equal to the sea leg's undegraded deliverable capacity. A state that has invested more heavily in its sea-based force therefore faces steeper penalty per unit of adversary improvement, not a shallower one.

Three mutually exclusive and exhaustive cases follow. A state is cliff-immune when the residual term meets or exceeds the floor, so sea-leg attrition cannot break the condition at any neutralisation probability. It is cliff-exposed when the residual term sits below the floor but above the floor less the sea leg's capacity, placing the cliff inside the probability interval as a genuine threshold. It is structurally insufficient when the residual term drops below that lower bound, yielding a negative formal solution — a statement about parameters, not posture. The deterrence condition cannot be satisfied even at zero neutralisation probability; Section 3.6 treats this as a corner solution, not as evidence that a state occupies a permanent reversion zone.

The comparative statics are direct. Deeper residual capability pushes the cliff outward at a rate given by the reciprocal of sea-leg capacity; more sea-based warheads also push it outward, though with diminishing returns. The requirement rises one-for-one with the assumed floor, which is why multiple floor values are carried throughout.

Two properties deserve explicit statement. First, the discontinuity is imposed rather than discovered: expected retaliatory capability declines continuously and linearly, and the cliff's abrupt character arises entirely from the binary threshold rule — a deliberate representation of assured destruction as a categorical political judgement. Readers who prefer a continuous-degradation interpretation may treat the cliff as a crossing point and the requirement as a continuously rising de-

mand; the principal instrument accommodates either reading without modification. Second, independence across warheads is an approximation. The expression treats neutralisation as independent across warheads, whereas in practice it is correlated at platform level, since warheads aboard one submarine are neutralised jointly and boats sharing patrol geography face correlated conditions. The consequences are quantified in Section 3.4. Decoys, false contacts, multiple patrol areas, differing platform signatures, and environmental variability all reduce the adversary's chain-completion probability and are absorbed as downward pressure on p rather than modelled separately.

3.3. Parameters and the Derivation of the Residual Retaliation Term

The model draws on two kinds of parameter. Source-derived quantities such as force structures, missile ranges, platform configurations, and computed distances, are reported in full, with their evidence categories and sources, in Table A1 in the Appendix, kept separate from author-assigned parameters below in accordance with the distinction the analysis maintains throughout. Two entries in that table warrant emphasis here because they bind results that follow: every quantity carries one of seven evidence categories defined in Section 2.3 rather than being described simply as verified, and Arihant-class launch-tube configuration is the constraint which converts the platform counts into deliverable warheads, it makes the sea-based figures below substantially smaller than raw warhead totals would imply.

Table 1 reports the author-assigned parameters, each stated as an interval rather than a point value and each sampled across its full width in the simulations of Section 3.4, so that no conclusion rests on a single assignment (Table 1).

Table 1. Author-assigned parameters and rationale.

Parameter	Range	Basis for the assigned range
N sea, Pakistan	8 to 24 warheads	Agosta-90B force; Babur-3 loadout assumed
N sea, India, China mission	4 to 8 warheads	One or two boats at 4 K-4 each
N sea, China	72 to 144 warheads	6 Type-094/094A boats, 12 tubes each
P(D) sea, Pakistan	0.35 to 0.55	Reach binding; limited test record
P(D) sea, India, China mission	0.75 to 0.85	Reach satisfied; K-15 and K-4 record
P(D) sea, China	0.78 to 0.85	Reach satisfied; mature systems
Land-leg survival	0.05 to 0.90	Four scenarios; see Table 2
P(D) land	0.64 to 0.855	Reliability by penetration
Deterrence floor	5, 10, 20	5 sourced; 10 and 20 conservative
Belief standard deviation	0.05 to 0.15	Opacity of capability indicators
Utility ratio	0.05 to 0.30	Illustrative calibration
Hedge requirement	1 to 5 warheads	Minimum buy-down justifying hedge

The most authoritative treatment of the parameter on which the residual term turns is the analysis of Indian counterforce capability by Clary and Narang (2019). It yields three anchors. On force size, it describes Pakistan's several dozen strategic land assets and poses a disarming hypothetical in which destruction of ninety per cent would leave approximately five strategic missiles, implying roughly fifty India-capable strategic warheads, consistent with Kristensen et al. (2023) net of tactical, air-delivered, and reserve allocations. On deterrence floor, it states that those residual five would be sufficient to deter any reasonable Indian leader, providing dyad-specific peer-reviewed evidence for a minimum floor; the value of five is therefore adopted as sourced lower bound, with ten and twenty retained as conservative alternatives. On survivability, its own verdict is that Indian counterforce

prospects are dubious, documenting inexpensive decoys, hardened underground storage functioning as a warhead sponge, routine dispersal, and the 1991 precedent in which approximately 1,500 strikes destroyed likely no mobile launchers.

Because the literature does not support a single survivability value, land-leg survival is carried as four labelled scenarios, and every subsequent result is reported against all four. (Table 2).

Table 2. Land-leg survivability scenarios and the resulting residual retaliation term for Pakistan.

Scenario	Survival	Basis	*A* (5th / 50th / 95th)
S1 Highly effective	0.05 to 0.15	90 per cent destruction case	2.0 / 3.7 / 5.8
S2 Moderately effective	0.15 to 0.35	Fixed sites hit, mobiles survive	5.6 / 9.1 / 13.7
S3 Dubious (sourced central)	0.30 to 0.60	Decoys, sponge sites, dispersal	10.9 / 16.5 / 23.6
S4 Mobile-hunt analogue	0.60 to 0.90	1991 precedent	20.4 / 27.5 / 36.5

3.4. Sensitivity Analysis

An earlier version of this analysis reported a hand-calculated sensitivity grid containing an arithmetic error, describing seventeen parameter combinations as eighteen, and examined only selected combinations. Both defects are corrected by replacing selective enumeration with complete enumeration and probabilistic propagation.

Neutralisation is simulated at platform level using a one-factor Gaussian copula, in which a boat is neutralised when a latent variable combining a patrol-area common factor and an independent idiosyncratic factor falls below the threshold corresponding to the marginal neutralisation probability. This preserves the marginal probability exactly while inducing correlated outcomes. The aggregation and correlation configurations are summarized below (Table 3).

Table 3. Effect of platform aggregation and correlation on surviving warheads (16 warheads, neutralisation probability 0.4).

Model	Mean	SD	P(S ≤ 2)	P(S = 0)
Warhead-level, independent	9.60	1.96	0.0 %	0.0 %
2 boats × 8, ρ = 0	9.60	5.54	16.0 %	16.0 %
2 boats × 8, ρ = 0.5	9.59	6.39	24.0 %	24.0 %
4 boats × 4, ρ = 0	9.60	3.92	2.6 %	2.6 %
4 boats × 4, ρ = 0.5	9.61	5.52	12.6 %	12.6 %

The expectation is unchanged in every configuration, confirming that point estimates are unbiased under aggregation. Variance rises with aggregation and correlation, and the direction of resulting effect on deterrence condition depends on margin: higher variance increases probability of falling below the floor when the mean lies above it, and decreases that probability when the mean lies below it. A state operating with a thin surplus is therefore harmed by force concentration, while a state already in deficit is helped by it. A policy-relevant implication follows: distributing same sixteen warheads across four boats rather than two reduces probability of total loss from 16.0 to 2.6 per cent, so dispersal functions as a survivability measure independent of any change in adversary capability. The corresponding Monte Carlo classification for India's China-mission leg is reported separately (Table 4).

Table 4. Monte Carlo classification of India's China-mission leg (200,000 draws).
Residual term: median 7.6, 95 per cent credible interval 3.3 to 14.8. Sea-based capacity: median 4.8 warheads.

Floor	Immune	Exposed	Insuff.	Median p^*	95 % CrI	Median p_r
5	82.8 %	17.2 %	0.0 %	0.82	0.45 to 0.99	0.72
10	25.3 %	54.2 %	20.5 %	0.45	0.03 to 0.97	0.35
20	0.0 %	2.1 %	97.9 %	0.15	0.01 to 0.55	0.06

Two features of India result are worth isolating before comparison. The classification is strongly sensitive to the assumed deterrence floor: at the sourced floor of five, India is cliff-immune in more than four draws out of five, whereas at conservative floor of twenty it is structurally insufficient in almost all of them. In intermediate case, which the analysis treats as its reference, threshold falls near middle of the probability interval, and its credible interval is wide enough that the cliff could lie anywhere between near-total neutralisation and a level of capability already claimed in published simulation. Pakistan cannot be classified in the same single table, because its residual retaliation term is precisely the quantity the literature disputes; it is therefore reported against all four survivability scenarios. The corresponding Pakistan classifications are reported separately (Table 5).

Table 5. Monte Carlo classification of Pakistan by survivability scenario (200,000 draws).

Scenario	Floor = 5	Floor = 10	Floor = 20
S1	Immune 16 %; exposed 84 %	Exposed 61 %; insufficient 39 %	Insufficient 100 %
S2	Immune 99 %	Immune 39 %; exposed 60 %	Insufficient 84 %
S3 (sourced central)	Immune 100 %	Immune 98 %	Immune 22 %; exposed 58 %
S4	Immune 100 %	Immune 100 %	Immune 96 %

China is cliff-immune in 100 per cent of draws at every floor tested. The anticipation margin across all propagated parameters has a median of 0.089 and a 95 per cent credible interval of 0.037 to 0.198. A complete twenty-seven-cell deterministic grid is reported in Appendix A.5.

3.5. The Posture Reversion Mechanism

A state holds beliefs about the adversary's current aggregated neutralisation capability, represented as a random variable on the bounded unit interval with a mean and a standard deviation. Let the operative structural threshold be the credibility cliff for a cliff-exposed state, or the hedge threshold derived below. Each period the state chooses between assured retaliation and asymmetric escalation, with von Neumann–Morgenstern utilities on a common arbitrary scale: V is the deterrence dividend from maintaining assured retaliation while the condition holds; $-C$ is the cost of maintaining that posture when the condition has in fact failed and the posture is exposed as hollow; and R_0 is the certain payoff from reverting. The ordering is that V exceeds R_0, which exceeds $-C$. No identification of the catastrophic cost with a warhead count is made; the earlier setting of that cost equal to the deterrence floor equated a utility with a force quantity and has been removed.

Let q denote the subjective probability that the threshold has already been crossed. Expected utility from assured retaliation is $V - q(V + C)$, and from reversion is R_0. The state reverts exactly when the former falls below the latter, that is, when

$$q > (V - R_0) / (V + C) \equiv \alpha$$

This is a probability threshold, and α is dimensionless, being a ratio of utility differences, so that the arbitrary scale cancels and never enters a reported number. Substantively, α is the probability of the threshold-crossing that a state tolerates before abandoning assured retaliation: a state regarding exposure as catastrophic, or holding an attractive fallback, has a small α and reverts on the weaker evidence.

Defining the reversion threshold as the belief mean at which the condition binds, under the normal beliefs the threshold is the structural threshold less the product of standard normal quantile at one minus α and the belief standard deviation. This is exact rather than approximate; no Taylor expansion is required, and earlier expression, which set anticipation margin proportional to deterrence floor divided by deterrence dividend, does not follow from stated utilities and is withdrawn.

A normal distribution assigns positive probability outside unit interval, that is inadmissible for a belief about probability. The beta family is therefore adopted as baseline, matched on mean and standard deviation, and all reported thresholds are computed numerically from it. The normal closed form is retained as interpretive device with its accuracy quantified: across operating region the maximum discrepancy between beta and normal reversion thresholds is 0.052 and median is below 0.005, degrading only at high belief means with wide uncertainty.

The following proposition replaces the earlier appeal to unnamed variance-dominance property, for which no supporting theorem was available.

Proposition. Under normal beliefs, the anticipation margin is strictly increasing in belief uncertainty if and only if α is less than one half, which holds if and only if catastrophic cost exceeds the deterrence dividend less twice the reversion fallback.

Proof. The margin is linear in belief standard deviation with slope equal to standard normal quantile at one minus α, which is strictly positive if and only if α is less than one half. Substituting definition of α and rearranging gives the stated condition. If catastrophic cost is at least the deterrence dividend and the fallback is non-negative, the condition holds automatically. Under beta beliefs the same monotonicity is verified numerically across full calibration range. The substantive reading is unchanged but now rests on a proved statement: opacity is itself destabilising, since a state unable to observe adversary capability precisely reverts earlier than one that can, holding capability fixed.

For a cliff-exposed state the mechanism applies as derived. For a structurally insufficient state the probability of threshold-crossing equals one for every admissible belief, so the expected utility of assured retaliation is always below that of reversion: the posture is never rationalisable and there is no reversion interval, because there is nothing to revert from. The observable implication is a stable corner solution, which is Proposition P3. No infeasible thresholds are reported anywhere in this analysis.

A cliff-immune state may nonetheless maintain a sea-based leg as insurance against deterioration elsewhere. Define the hedge value as the difference between the residual requirement at complete neutralisation and the requirement at the current neutralisation probability, which equals the lesser of the deterrence floor and the sea leg's remaining deliverable capacity, and which declines linearly to zero. With the hedge requirement denoting the minimum buy-down justifying the programme, expressed in warhead units and therefore commensurate with the hedge value, the hedge is foreclosed when the hedge value falls below the requirement, giving a hedge threshold equal to one less the ratio of the requirement to the sea leg's capacity. The anticipation structure is identical, with the hedge threshold substituting for the cliff. One mechanism therefore covers two threshold definitions: the credibility cliff for states whose deterrence condition can break, and hedge foreclosure for states whose insurance can be devalued. The conceptual sequence is shown schematically (Figure 1).

Faculty of Security Studies
University of Belgrade

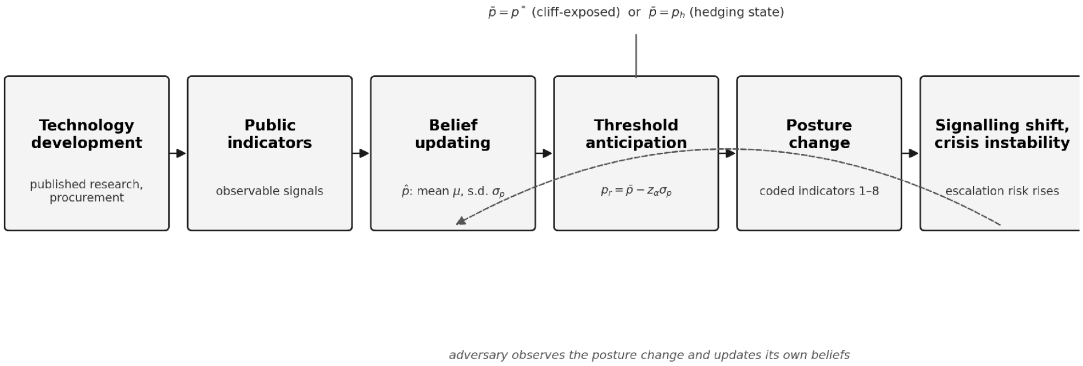


Figure 1. The conceptual sequence formalised in this article, from capability development through public indicators and belief updating to threshold anticipation, posture change, and crisis instability. The operative threshold is the credibility cliff for a cliff-exposed state and the hedge-foreclosure threshold for a cliff-immune state maintaining a sea-based hedge. The dashed path indicates that an adversary observing the posture change updates its own beliefs, so that the sequence is recursive rather than terminal.

3.6. The Strategic Interaction

The setting is a two-player normal-form game between China, choosing between deployment of AI-enabled anti-submarine warfare at scale and restraint, and India, choosing between expansion of its sea-based force with its protective screen and restraint. Pakistan holds no strategy set and bears an externality. This is stated plainly because the earlier presentation as a three-player game was mathematically incorrect: a player without a strategy set is not a player. Endowing Pakistan with strategies would misrepresent the structural finding, since within model's horizon Pakistan holds no posture instrument that alters either player's payoffs, and the externality formulation makes that asymmetry visible. The externality is specified as hedge foreclosure rather than deterrence collapse.

Payoff orderings are author assignments, which are justified cell by cell. For China, deployment against Indian restraint ranks highest, delivering capability against United States and Indian sea-based forces with no offsetting expansion; deployment against Indian expansion ranks second, the same capability with regional gain partially offset; mutual restraint ranks third; and restraint against Indian expansion ranks lowest. For India, mutual restraint ranks highest, with no threat and no expansion cost; Chinese restraint with Indian expansion ranks second, imposing unnecessary cost; Chinese deployment with Indian expansion ranks third, the threat partially offset; and Chinese deployment with Indian restraint ranks lowest, the margin conceded.

The dominance result rests on explicit premise, Assumption G1: the value to China of capability against United States sea-based forces exceeds the costs China attaches to regional instability, arms racing, and induced countermeasures. This is an assumption about preferences, supported qualitatively by scale and mission scope of documented Chinese investment but not demonstrated by force-structure data. Rather than defend it, its role is quantified by complete ordinal enumeration. Each player has twenty-four strict orderings over four outcomes, giving 576 ordinal games. The deployment-expansion outcome is a Nash equilibrium in 144 (25.0 per cent) and the unique pure-strategy equilibrium in 108 (18.8 per cent). Imposing four substantive constraints encoded in cell justifications narrows the space to six admissible profiles: where Assumption G1 holds, the deployment-expansion outcome is the unique equilibrium in both applicable profiles; where it fails, mutual restraint is an equilibrium in three of four remaining profiles, and the interaction becomes a coordination problem rather than a dominance trap. The resulting ordinal payoffs are summarized below (Table 6).

Table 6. Strategic interaction: ordinal payoffs (China, India). Pakistan bears the externality at its worst value under deployment-expansion outcome and on the capability channel wherever China deploys.

Payoffs (China, India)	India expands	India restrains
China deploys	(3, 2) equilibrium under G1	(4, 1)
China restrains	(1, 3)	(2, 4) mutual restraint

The outcome of mutual restraint is not the cooperative equilibrium. Under stated orderings China obtains three at equilibrium and two at mutual restraint, and therefore prefers the equilibrium; India prefers mutual restraint. Mutual restraint is thus not a jointly preferred outcome blocked by strategic structure, but one player's preference set against another's, which is a conflict of interest rather than collective-action failure. The earlier label is withdrawn on the both counts.

Three qualifications discipline the one-shot result. By folk-theorem logic (Fudenberg & Tirole, 1991), mutual restraint could be sustained in indefinitely repeated game with sufficiently patient players and observable defection; the binding constraint is observability, since AI-enabled capability resides in software, the training data, and fusion architectures without inspectable physical signatures, and issue linkage and side payments inherit same constraint. Because Assumption G1 makes the United States force posture the principal driver of Chinese incentives, regional game is properly understood as nested within the wider strategic interaction (Tsebelis, 1990), so that regional equilibrium may shift through changes in arena the regional players do not control. What the model supports is accordingly bounded: in its one-shot form, under Assumption G1 and four documented constraints, no bilateral agreement among regional states alters the equilibrium.

3.7. Coding Framework for Observable Posture Change

Propositions P1 and P3 require observable indicators. Eight are specified: alert status; deployment pattern; nuclear signalling; launch authority; first-use doctrine; patrol behaviour; force dispersal; and countermeasure investment. Coding requires that a change be documented in at least one source of the stated type, dated, and assessed against the evidence categories in Section 2.3. Indicator movement alone does not establish the mechanism; attribution to beliefs about anti-submarine warfare capability requires either explicit official reasoning or the exclusion of competing explanations, and neither is available at the unclassified level for any case examined here.

3.8. Case Results

Under the sourced central scenario and the sourced deterrence floor, Pakistan is cliff-immune in essentially all draws. Its assured retaliation does not rest on the sea leg and is not abolished by AI-enabled anti-submarine warfare. Sustaining the characterisation advanced in an earlier version of this framework would require simultaneously adopting the most India-favourable survivability assumption and the most demanding floor, a combination no source supports and the governing literature contradicts. That claim is therefore withdrawn.

At sixteen warheads on patrol and a delivery probability of 0.45, complete neutralisation of the sea leg raises the required residual survivability by 5.0 warheads at a floor of five and by 7.2 warheads at a floor of ten, which is approximately ten to fourteen percentage points of a fifty-warhead strategic force. Where Pakistan requires the hedge to buy down at least two warheads of requirement, foreclosure occurs at a hedge threshold of 0.72, with a 95 per cent credible interval of 0.43 to 0.83, and anticipation at 0.61.

India's China-mission leg rests on a materially thinner sea-based force than warhead totals alone suggest. The Arihant-class boats carry four launch tubes, and each tube accommodates either three K-15 or a single K-4 (Overt Defense, 2024). Because the K-15's 750-kilometre range cannot reach

Chinese targets from Bay of Bengal, only K-4 rounds count toward this mission, giving a maximum of four per boat and eight across two boats commissioned as of the evidence cut-off. The third boat, carrying eight tubes, entered service after that date and is excluded.

On that basis, India's China-mission leg is cliff-exposed in 54.2 per cent of draws at floor of ten, with a median cliff of 0.45 and median reversion threshold of 0.35; at central calibration the danger interval runs from 0.34 to 0.45. This is case in which the credibility cliff is operative concept and the mechanism applies with an interior rather than corner threshold. Two features distinguish it from Pakistani case. First, the threshold sits near middle of the probability interval rather than at its extremes, so anticipatory zone lies within range of the capability improvement currently under way rather than beyond it. Second, India is structurally insufficient in a fifth of draws at same floor and in almost all draws at the conservative floor of twenty: where the deterrence requirement is set high, four to eight deliverable warheads cannot meet it whatever the adversary's capability. The mission-specific qualification stated in Section 1.3 is therefore doing substantive work here, since this is a claim about one leg against one adversary and not about Indian retaliatory capability as a whole.

An earlier version presented Indian patrol scheduling as empirical confirmation of mechanism. No specific verifiable evidence supports that attribution, and the claim is withdrawn as stated. Indian patrol practice, investment in anti-submarine capability in the Bay of Bengal, and public commentary on the bastion defence (Indian Defence News, 2025; Swarajya, 2025) are consistent with prediction, but equally consistent with the routine force maturation following commissioning, with the conventional naval competition, and with organisational momentum. Chinese research and survey vessel deployments in Indian Ocean, which bear directly on the oceanographic data that fusion-based detection requires, are documented in open reporting (Indo-Pacific Defense Forum, 2026); no public source, however, links them to any specific Indian posture decision.

China is cliff-immune in every draw at every floor. China's sea leg is one element of diversified force whose residual legs alone clear any floor considered here. What matters about China in this system is not vulnerability but pressure: under Assumption G1, Chinese ASW development is driven by competition with United States and it produces regional consequences neither India nor Pakistan can deflect.

If read together, the three cases show that the framework generates differentiated, falsifiable classifications when calibrated from declared sources. What they do not show is that any state has already changed posture because it judged adversary capability to be nearing a threshold, no unclassified record establishes how these leaderships actually assess AI-enabled ASW or weight the utilities entering reversion condition. The propositions remain predictions with specified disconfirming evidence, not confirmed findings; the empirical contribution is to demonstrate that the mechanism is operative and measurable, not that it has been observed.

4. Discussion

4.1. Interpretation in Relation to Prior Research

The results qualify rather than overturn the counterforce literature. Lieber and Press (2017) argue that technological change has moved the balance toward counterforce feasibility; the present analysis is consistent with that direction but locates the effect differently. In the asymmetric dyad, emerging sensing capability does not first threaten deterrent as a whole; it first devalues hedge that a state is acquiring against erosion of its principal leg. This is narrower mechanism than deterrence collapse and operates earlier.

The finding engages Clary and Narang (2019) directly. Their assessment that Indian counterforce prospects are dubious places Pakistan in cliff-immune category under central scenario, and their observation that anticipation of Indian counterforce development would accelerate Pakistan's shift toward a sea-based forces supplies the mechanism by which hedge acquires its value. The present contribution renders that hedge measurable and shows it can be foreclosed by adversary capability who never touches the land-based force.

With respect to Long and Green (2015), analysis supports their emphasis on intelligence and sensing over the warhead counts, it adds that strategically decisive quantity may be the survivability burden transferred between legs rather than the probability of successful counterforce as such. With respect to Narang (2014), the corner-solution result offers formal account of why asymmetric-escalation postures may prove insensitive to adversary capability milestones: for structurally insufficient state, assured retaliation is not a foregone option but unavailable one.

Geography and force structure constrain state's option set; they do not determine doctrine. Posture also reflects political preferences, civil–military relations, threat perception, organisational interests, and historical experience, and strategic utility may incorporate identity, legitimacy, and doctrinal value rather than the material capability alone (Hetteh, 2026). This is load-bearing in present case: because Pakistan's sea-based programme is not fully rationalisable through assured retaliation, non-material components carry the residual explanation — ambiguity imposed on adversary planning, status, and doctrinal legitimacy.

4.2. Arms Control Implications

The following are conditional consequences of model rather than unconditional prescriptions, and each carries a feasibility constraint that must be stated alongside it.

Oceanographic transparency and the survey data-sharing are conditional on the Assumption G1 failing, or on repeated interaction. The constraint is that scientific data are dual-use, so a regime must protect the legitimate research while resisting the strategic deception, since degraded or misleading environmental data can be shared at low cost.

Patrol-area or anti-submarine restraint understandings are conditional on cliff-exposed case. The constraint is secure-reserve dilemma: declaring a bastion protects a deterrent only if the declaration is not itself exploitable targeting information, and reciprocity is difficult to confirm when the restrained capability is algorithmic.

Trilateral risk-reduction dialogue is conditional on two players recognising externality. The constraint is that participation may itself reveal capability levels through act of specifying what is restrained.

Algorithmic restraint measures face binding constraint that no verification instrument exists for a capability without physical signature. Every prior regime, whether limiting defensive systems, intermediate-range missiles, or strategic launchers, counted or observed physical objects. This is presented as research agenda rather than adoptable proposal.

4.3. Artificial Intelligence Failure Modes

AI-enabled anti-submarine warfare inherits the failure modes of large-scale autonomous inference: false positives against a background of biologics and decoys, error propagation across fused sensor feeds, disagreement between modalities, vulnerability to data poisoning and the adversarial deception, and compression of human decision time precisely where classification errors are most consequential. A false positive that persuades an adversary its deterrent has been compromised is destabilising even if underlying capability does not exist, a mechanism operating through same belief channel modelled in Section 3.5 and one that strengthens rather than weakens the argument. Governance of such systems cannot therefore be reduced to counting rules, and belongs to emerging regulatory agenda for dual-use autonomous technologies, where deployment currently outpaces ethical and legal frameworks (Beriša, Cvetković, & Pavić, 2024; Ivanov, Babanoski, & Cvetković, 2026).

Faculty of Security Studies
University of Belgrade

4.4. Responsible Publication and Dual-Use Assessment

All inputs to this analysis are public: published force estimates, officially stated missile ranges, city coordinates from an open gazetteer, and open reporting on survey activity. All derived quantities are aggregated to country-level parameter ranges. Patrol geography is characterised only at the level already present in the cited open literature. No operating areas, schedules, platform-specific vulnerabilities, or targeting-relevant parameters beyond the public record are derived or disclosed, and countermeasures are discussed only as categories, without operational parameters. The analysis was reviewed against the standard that no combination of its aggregated inputs yields targeting-relevant information absent from its sources individually.

4.5. Limitations

Six limitations bound the analysis. The additive force representation, the binary threshold rule, the one-shot game, and the correlation structure are specifications rather than discoveries, and the threshold discontinuity is imposed by the deterrence rule rather than present in the capability function. Delivery-probability components, survivability scenarios, deterrence-floor bounds, the hedge requirement, the utility ratio, the belief standard deviation, and the payoff orderings are author assignments governed by stated rules; all scoring is the work of a single author, no second independent coder or formal expert elicitation having been available, and while Monte Carlo propagation quantifies this uncertainty it does not remove it. The underlying force estimates are themselves analytic products, they carry unpublished uncertainty. There is no unclassified documentation exists on how any of three leaderships assesses capability trajectories or weights utilities entering reversion condition, so the belief and utility parameters are illustrative calibrations rather than estimates. The headline Chinese figures are simulation results reported through secondary coverage of study whose original-language version was not consulted, and gap between simulated and wartime performance is unmeasured. Finally, all force-structure, platform, and event claims are stated as of 31 December 2025.

Within these limits, framework's value is that every contestable input is explicit, bounded, and varied, so that disagreement can be conducted over declared parameters rather than hidden assumptions.

University of Belgrade

4.6. Future Research

Three extensions follow directly. A full three-player model incorporating the United States as a strategic actor would test whether the nested structure identified in Section 3.6 alters the regional equilibrium, and would connect the framework to the extended-deterrence and alliance-coordination literature. Independent coding or formal expert elicitation of the survivability and delivery parameters would convert the present single-author scoring into a validated measurement instrument. Application to other asymmetric dyads with sea-based hedges, for which the scope conditions in Section 1.3 hold, would test Proposition P2 across cases rather than within one.

5. Conclusions

This article has developed two dual formal instruments for analysing anticipatory posture change under emerging counterforce technology, and applied them to AI-enabled anti-submarine warfare in the China–India–Pakistan triangle.

Mathematically derived are the three-case classification of deterrence conditions; the residual retaliation requirement and its comparative statics, including the result that survivability burden transfers to residual legs at a rate equal to the sea leg's undegraded deliverable capacity; the reversion condition as a probability threshold; the closed-form anticipation margin, exact under normal

beliefs; the monotonicity proposition and its stated condition; the corner-solution result for structurally insufficient states; the hedge-foreclosure threshold; and the equilibrium of the two-player interaction conditional on Assumption G1.

Empirically observed, with sources categorised and evidence cut-off of 31 December 2025, are force structures, platform commissioning, missile tests, published simulation claims, survey activity, and declaratory doctrine. Remaining as theoretical prediction rather than demonstrated finding are Propositions P1, P2, and P3. These categories are not conflated, and nothing in second confirms the third.

The central substantive claim is modest and, for that reason, defensible: AI-enabled anti-submarine warfare does not abolish assured retaliation in asymmetric dyads. It transfers survivability burden onto residual legs at a calculable rate, and it devalues sea-based hedges before it disables sea-based deterrents. This is why the behaviour it induces appears before the capability on which it depends can be verified, and it is why analysis of such technologies should attend to hedges and thresholds rather than to collapse.

The practical recommendation that follows most directly is also the least demanding: because dispersal reduces the probability of total loss independently of any change in adversary capability, and because verification instruments for algorithmic capability do not presently exist, force dispersal and the development of verification concepts for capabilities without physical signature deserve priority over restraint proposals that current instruments cannot support.

Author Contributions: The author is solely responsible for conceptualisation, methodology, formal analysis, investigation, writing of the original draft, and writing, review and editing.

Funding: This research received no external funding.

Acknowledgments: The author thanks the reviewers of this journal, whose detailed comments on an earlier version led to the reconstruction of the formal model and the withdrawal of a finding that could not be sustained.

Conflicts of Interest: The author declares no conflict of interest.

Data and Code Availability: All data derive from the published open sources cited. Complete simulation code, the random seed, coordinate data, distance calculations, and deterministic grids are provided in Appendix A, sufficient for full replication of every quantitative result.

Potentially Sensitive Open-Source Military Information: The manuscript uses only publicly available military information, aggregated as described in Section 4.4; no operationally sensitive detail beyond cited public record is derived or disclosed.

Appendix A

A.1. Evidence cut-off. All force-structure, platform, and event claims are stated as of 31 December 2025. The cut-off governs which events are included, not the publication dates of the sources reporting them; where a later publication documents an event falling before the cut-off, it is cited for that event only.

A.2. Model equations. Expected retaliation: $\mathrm{E}\underline{R} = N_{\mathrm{sea}}(1-p)P(\mathrm{D})_{\mathrm{sea}} + A$. Deterrence condition: $\mathrm{E}\underline{R} \geq D_{\min}$. Credibility cliff: $p^* = 1 - (D_{\min} - A)/(N_{\mathrm{sea}}P(\mathrm{D})_{\mathrm{sea}})$. Residual retaliation requirement: $A^*(p) = \max\{0, D_{\min} - N_{\mathrm{sea}}(1-p)P(\mathrm{D})_{\mathrm{sea}}\}$. Hedge value: $H(p) = \min\{D_{\min}, N_{\mathrm{sea}}P(\mathrm{D})_{\mathrm{sea}}(1-p)\}$. Hedge threshold: $p_{\mathrm{h}} = 1 - \kappa/(N_{\mathrm{sea}}P(\mathrm{D})_{\mathrm{sea}})$. Reversion condition: revert if $\Pr(\hat{p} > p^{-}) > \alpha = (V - R_0)/(V + C)$. Anticipation margin: $\delta = z_\alpha \sigma_{\mathrm{p}}$, with reversion threshold $p_{\mathrm{r}} = p^{-} - \delta$. Residual term: $A = \Sigma N_j s_j P(\mathrm{D})_j$ over land and air legs. Delivery decomposition: $P(\mathrm{D})$ = reach × reliability × penetration.

A.3. Coordinates and distances. Source: GEOnet Names Server. Haversine formula, Earth radius 6,371 km. Karachi (24.8607° N, 67.0011° E); Mumbai (19.0760° N, 72.8777° E), 883.4 km; Ahmedabad (23.0225° N, 72.5714° E), 601.8 km; New Delhi (28.6139° N, 77.2090° E), 1,095.7 km. At the stated 450 km range, reach requires forward deployment of 433.4 km for Mumbai, 151.8 km for Ahmedabad, and 645.7 km for New Delhi.

A.4. Source-derived parameters. Table A1 reports every quantity taken from a source rather than assigned by the author, each with its evidence category (Section 2.3) and origin. It is the source-derived counterpart to the author-assigned parameters in Table 1. (Table A1).

Table A1. Source-derived parameters, with evidence categories.

Parameter	Value or range	Evidence category	Source
Pakistan sea platforms	3 Agosta-90B with AIP	Officially reported	Kristensen et al. (2023)
Pakistan Hangor-class	8 on order; launches from 2024	Independently assessed	Defence Security Asia (2025)
Pakistan total warheads	approx. 170	Independently assessed	Kristensen et al. (2023)
Babur-3 SLCM range	450 km (stated)	Officially reported	ISPR (2017); Panda (2018)
Pakistan strategic land force	approx. 50 India-capable warheads	Estimated	Clary and Narang (2019)
India SSBN force	Arihant 2016; Arighat 2024	Confirmed	Kristensen et al. (2024)
Arihant-class SLBM tubes	4 tubes; 4 K-4 or 12 K-15	Independently assessed	Overt Defense (2024)
K-4 SLBM range	approx. 3,500 km	Officially reported	Kristensen et al. (2024)
K-4 first submarine test	27 November 2024	Confirmed	Naval Technology (2024)
China SSBN force	6 Type-094/094A	Independently assessed	U.S. Department of Defense (2024)
JL-2 and JL-3 range	approx. 7,200 and 10,000 km	Projected	U.S. Department of Defense (2024)
China warheads	more than 600	Independently assessed	U.S. Department of Defense (2024)
Target distances	602 to 1,096 km from Karachi	Confirmed (computed)	Appendix A.3

Faculty of Security Studies
University of Belgrade

A.5. Simulation procedure. 200,000 draws per configuration; NumPy default generator, seed 20260722; uniform independent sampling across the ranges in Table 1; deterrence floors reported separately at 5, 10, and 20; platform-level neutralisation by one-factor Gaussian copula; reversion thresholds computed from the beta family, with the normal closed form used within the simulation and its maximum discrepancy of 0.052 in the operating region reported in Section 3.5.

A.6. Deterministic grid. The complete twenty-seven-cell enumeration underlying Section 3.4 is reported below (Table A2).

Table A2. Deterministic grid of the Pakistan hedge threshold. Entries are hedge thresholds; F denotes a hedge already foreclosed at zero neutralisation probability.

N sea	P(D) sea	$\kappa = 1$	$\kappa = 2$	$\kappa = 3$
8	0.35	0.64	0.29	F
8	0.45	0.72	0.44	0.17
8	0.55	0.77	0.55	0.32
16	0.35	0.82	0.64	0.46
16	0.45	0.86	0.72	0.58
16	0.55	0.89	0.77	0.66
24	0.35	0.88	0.76	0.64
24	0.45	0.91	0.81	0.72
24	0.55	0.92	0.85	0.77

A.7. Code.

```
    import numpy as np
from scipy import stats, optimize
rng = np.random.default_rng(20260722); M = 200_000

def classify(N_rng, P_rng, A_rng, D_min):
    N = rng.uniform(*N_rng, M); P = rng.uniform(*P_rng, M); A = rng.uniform(*A_rng, M)
    need, cap = D_min - A, N * P
    immune = need <= 0
    insuff = (~immune) & (cap < need)
    exposed = (~immune) & (~insuff)
    pstar = np.where(exposed, 1 - need / cap, np.nan)
    return immune.mean(), exposed.mean(), insuff.mean(), pstar

alpha = rng.uniform(0.05, 0.30, M); sigma = rng.uniform(0.05, 0.15, M)
delta = stats.norm.ppf(1 - alpha) * sigma

def hedge(N_rng, P_rng, kappa):
    N = rng.uniform(*N_rng, M); P = rng.uniform(*P_rng, M)
    return 1 - kappa / (N * P)

def surviving(n_boats, w, p, rho):
    W = rng.standard_normal((M, 1)); eps = rng.standard_normal((M, n_boats))
    Z = np.sqrt(rho) * W + np.sqrt(1 - rho) * eps
    return (n_boats - (Z < stats.norm.ppf(p)).sum(1)) * w

def pr_beta(pbar, s, a):
    def q(mu):
        k = mu * (1 - mu) / s**2 - 1
        return 1 - stats.beta.cdf(pbar, mu * k, (1 - mu) * k)
    lo = (1 - np.sqrt(1 - 4 * 1.05 * s**2)) / 2 + 1e-6
    return optimize.brentq(lambda m: q(m) - a, lo, pbar - 1e-6, xtol=1e-12)

def hav(lat1, lon1, lat2, lon2, R=6371.0):
    p1, p2 = np.radians(lat1), np.radians(lat2)
    dp, dl = p2 - p1, np.radians(lon2 - lon1)
    h = np.sin(dp/2)**2 + np.cos(p1)*np.cos(p2)*np.sin(dl/2)**2
    return 2 * R * np.arcsin(np.sqrt(h))
```

Faculty of Security Studies
University of Belgrade

Appendix B

B.1. Baseline, Assumption G1 holding. China: deploy-restrain 4, deploy-expand 3, restrain-restrain 2, restrain-expand 1. India: restrain-restrain 4, restrain-expand 3, deploy-expand 2, deploy-restrain 1. Deployment is strictly dominant for China; the unique pure-strategy equilibrium is deployment with expansion.

B.2. Variant 1, Assumption G1 failing. China: restrain-restrain 4, deploy-restrain 3, restrain-expand 2, deploy-expand 1; India unchanged. Deployment is no longer dominant; mutual restraint becomes an equilibrium and the interaction is a coordination problem.

B.3. Variant 2, Indian restraint under Chinese deployment. India prefers restraint to expansion when China deploys. The equilibrium becomes deployment with Indian restraint; the Pakistan externality is mitigated on the expansion channel but persists on the capability channel.

B.4. Complete enumeration. All 576 strict ordinal games were computed. Six admissible profiles survive the four documented substantive constraints, with equilibria as reported in Section 3.6.

B.5. What survives every variant. In all admissible profiles Pakistan's outcome is determined entirely by choices made by others. The specific deployment-expansion equilibrium survives only under Assumption G1 with baseline Indian preferences, which is the precise sense in which the strategic conclusions of this article are conditional.

6. References

1. Beriša, H., Cvetković, V., & Pavić, A. (2024). Implications of artificial intelligence and cyberspace on risk management capabilities. *International Journal of Disaster Risk Management, 6*(2), 279–295. https://doi.org/10.18485/ijdrm.2024.6.2.18
2. Clary, C., & Narang, V. (2019). India's counterforce temptations: Strategic dilemmas, doctrine, and capabilities. *International Security, 43*(3), 7–52. https://doi.org/10.1162/isec_a_00340
3. Cvetković, V. (2026). *Disaster risk management: Theory, concepts and methods*. Belgrade: Scientific-Professional Society for Disaster Risk Management.
4. Defence Security Asia. (2025, August 17). *Pakistan Navy launches third Hangor-class submarine PNS Mangro: Strategic game-changer against India's Kalvari fleet*. Retrieved from https://defencesecurityasia.com/en/pakistan-navy-hangor-submarine-pns-mangro-vs-india-kalvari/
5. Edmond, C. (2026). The balance of power in Ukraine and deterrence effectiveness against nuclear-armed Russia. *International Journal of Contemporary Security Studies, 2*(1), 1–16. https://doi.org/10.18485/fb_ijcss.2026.2.1.1
6. Friedman, N. (2006). *The Naval Institute guide to world naval weapons systems* (5th ed.). Annapolis, MD: Naval Institute Press.
7. Fudenberg, D., & Tirole, J. (1991). *Game theory*. Cambridge, MA: MIT Press.
8. Hetteh, M. (2026). Revolutionary identity and nuclear bargaining: How symbolic legitimacy restructures utility in Iran's nuclear program. *International Journal of Contemporary Security Studies, 2*(1), 131–146. https://doi.org/10.18485/fb_ijcss.2026.2.1.11
9. Indian Defence News. (2025, December 1). *India's ASW imperatives for Bay of Bengal SSBN bastion defence*. Retrieved from https://www.indiandefensenews.in/2025/12/indias-asw-imperatives-for-bay-of.html
10. Indo-Pacific Defense Forum. (2026, February 13). *Missile test boosts India's nuclear submarine capability*. Retrieved from https://ipdefenseforum.com/2026/02/missile-test-boosts-indias-nuclear-submarine-capability/
11. International Institute for Strategic Studies. (2024). *The military balance 2024*. Abingdon, UK: Routledge.
12. Inter-Services Public Relations. (2017, January 9). *Press release: Test fire of SLCM Babur-3 (No. PR-11/2017-ISPR)*. Retrieved from https://ispr.gov.pk/press-release-detail.php?id=3570
13. Ivanov, A., Babanoski, K., & Cvetković, V. M. (2026). Iznad bojnog polja – etičke implikacije i regulatorni izazovi korišćenja autonomnih AI sistema za bezbednost životne sredine i zaštitu resursa (Beyond the battlefield: The ethical implications and regulatory challenges of using autonomous AI systems for environmental security and resource protection). *Serbian Journal of Engineering Management, 11*(1), 61–69. https://doi.org/10.5937/SJEM2601061I
14. Kristensen, H. M., Korda, M., & Johns, E. (2023). Pakistan nuclear weapons, 2023. *Bulletin of the Atomic Scientists, 79*(5), 329–345. https://doi.org/10.1080/00963402.2023.2245260
15. Kristensen, H. M., Korda, M., Johns, E., & Knight, M. (2024). Indian nuclear weapons, 2024. *Bulletin of the Atomic Scientists, 80*(5), 326–342. https://doi.org/10.1080/00963402.2024.2388470
16. Lieber, K. A., & Press, D. G. (2017). The new era of counterforce: Technological change and the future of nuclear deterrence. *International Security, 41*(4), 9–49. https://doi.org/10.1162/ISEC_a_00273

Faculty of Security Studies
University of Belgrade

17. Long, A., & Green, B. R. (2015). Stalking the secure second strike: Intelligence, counterforce, and nuclear strategy. *Journal of Strategic Studies, 38*(1–2), 38–73. https://doi.org/10.1080/01402390.2014.958150
18. Loss, R. (2019). *Artificial intelligence, the final piece to the counterforce puzzle? (LLNL-TR-791947). Livermore, CA: Lawrence Livermore National Laboratory*. Retrieved from https://www.osti.gov/biblio/1568008
19. Michell, S. (2021). AI in anti-submarine warfare (Interview with Rear Admiral René Tas). NITECH – NATO Innovation and Technology, (5), 60–62. Retrieved from https://issuu.com/globalmediapartners/docs/nitech_issue_05_jun_2021/60 Michell, S. (2021). AI in anti-submarine warfare (Interview with Rear Admiral René Tas). NITECH – NATO Innovation and Technology, (5), 60–62.
20. Mueller, J. (1988). The essential irrelevance of nuclear weapons: Stability in the postwar world. *International Security, 13*(2), 55–79. https://doi.org/10.2307/2538971
21. Mulyanyuma, A. A. (2026). United States–Republic of Korea joint nuclear defense guidelines: Evaluating alliance coordination against the North Korean threat. *International Journal of Contemporary Security Studies, 2*(1), 43–50. https://doi.org/10.18485/fb_ijcss.2026.2.1.4
22. Narang, V. (2014). *Nuclear strategy in the modern era: Regional powers and international conflict*. Princeton, NJ: Princeton University Press.
23. National Geospatial-Intelligence Agency. (n.d.). *GEOnet Names Server*. Retrieved from https://geonames.nga.mil
24. Naval Technology. (2024, November 29). *India tests nuclear-capable K-4 missile from INS Arighaat*. Retrieved from https://www.naval-technology.com/news/k-4-missile-from-ins-arighaat/
25. Overt Defense. (2024, August 29). *India commissions its second Arihant class SSBN – INS Arighat*. Retrieved from https://www.overtdefense.com/2024/08/29/india-commissions-its-second-arihant-class-ssbn-ins-arighat/
26. Panda, A. (2018, April 1). *Pakistan conducts second test of Babur-3 nuclear-capable submarine-launched cruise missile. The Diplomat*. Retrieved from https://thediplomat.com/2018/04/pakistan-conducts-second-test-of-babur-3-nuclear-capable-submarine-launched-cruise-missile/
27. Pavić, A. M., & Beriša, H. A. (2026). Multi-actor deterrence of small states: The case of the Republic of Serbia. *International Journal of Contemporary Security Studies, 2*(1), 29–42. https://doi.org/10.18485/fb_ijcss.2026.2.1.3
28. Powell, R. (2003). Nuclear deterrence theory, nuclear proliferation, and national missile defense. *International Security, 27*(4), 86–118. https://doi.org/10.1162/016228803321951108
29. Sagan, S. D. (1994). The perils of proliferation: Organization theory, deterrence theory, and the spread of nuclear weapons. *International Security, 18*(4), 66–107. https://doi.org/10.2307/2539178
30. South China Morning Post. (2025, September 14). *AI could cut a submarine's survival chance to 5%: Chinese defence scientists*. Retrieved from https://www.scmp.com/news/china/science/article/3325161/
31. Stefanick, T. (1987). *Strategic antisubmarine warfare and naval strategy*. Lexington, MA: Lexington Books.
32. Stupples, D. (2025, October 9). *AI, drone ships and new sensors could leave submarines with few places to hide in the ocean. The Conversation*. Retrieved from https://theconversation.com/ai-drone-ships-and-new-sensors-could-leave-submarines-with-few-places-to-hide-in-the-ocean-266973
33. Swarajya. (2025, December 25). *India's nuclear submarines can finally do what they were built for*. Retrieved from https://swarajyamag.com/defence/indias-nuclear-submarines-can-finally-do-what-they-were-built-for
34. Tsebelis, G. (1990). *Nested games: Rational choice in comparative politics*. Berkeley, CA: University of California Press. https://doi.org/10.1525/9780520911970
35. U.S. Department of Defense. (2024). Annual report to Congress: Military and security developments involving the People's Republic of China 2024. Washington, DC: Department of Defense.

36. Ultra Maritime. (2025, May 15). *Ultra Maritime showcases artificial intelligence for improved anti-submarine warfare mission effectiveness at the Combined Naval Event 2025 (Press release)*. Retrieved from https://umaritime.com/ultra-maritime-showcases-artificial-intelligence-cne-2025/
37. Wang, H., Meng, H., Shen, J., Zheng, L., Gu, J., & Sun, L. (2025). On intelligent decision support technology for shipborne helicopter submarine search. *Electronics Optics & Control, 32*(9), 74. Retrieved from https://www.researching.cn/articles/OJ3152e644d8992ceb
38. Wohlstetter, A. (1959). The delicate balance of terror. Foreign Affairs, 37(2), 211–234. https://doi.org/10.2307/20029345
39. Yu, Z., Zhu, Y., Zhang, W., Jing, K., Wang, S., Chen, C., . . . Du, J. (2025). Experimental demonstration of diamond quantum vector magnetometer for deep-sea applications. *National Science Review, 12*(4), nwae478. https://doi.org/10.1093/nsr/nwae478

Faculty of Security Studies
University of Belgrade